

CPUTM

COMPUTER POWER USER

Intel CTO
Pat Gelsinger
Talks About
His Company's
Future



* **Google**
The Search Continues

Result Page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [Next](#)

* **A Look At:**
FrontPage 2002
GoLive 6.0
Dreamweaver 4



* **Make Your Old
Hardware
Live On**

* **3GIO**
Replacing
The PCI Bus

From Our Columnists

The Rise Of MMP
Chipsets & Memory Standards
Sharpening My Teeth
The Quickening
Snow White & The Seven
Formats (+ 1)
The Nature Of Computer
Security: An Interdisciplinary
Approach
Web Design Workflow: Part II
netradio.txt
Hooked On Photonics

Hardware Reviews

HP Wireless Router
IBM Deskstar 120GXP
Apple iMac With
SuperDrive & More


June 2002
\$5.95 U.S. \$7.95 Canada



Breaking 30 GHz

**How Chip Makers
Are Changing Your PC**

room to burn



VINYL



DIGITAL VIDEO



DATA BACKUP



ORIGINAL MUSIC



DIGITAL PHOTO




MP3





Get organized, and take your files with you wherever you go. Burn all your data, video, photography, MP3s, vinyl, cassettes and whatever else you've got to CD with Easy CD Creator® 5 Platinum, from Roxio. The best selling CD burning software in the world. You can even create your own personalized jewel cases for each disc. Hit roxio.com to find out how. Mac® users check out Roxio's Toast® 5.

Now Windows® XP compatible.



NEXT-GEN



Spotlight

54

Evolution Or Revolution?

Your PC In 2003

58

Processors

The Next Little Thing

64


Video Cards

GPUs Of The Future

68

Buses & Chipsets

Coming Soon To A PC Near You



Frontside

- 6 What's Happening
- 12 Digital Economy
- 14 **The Saint**
The Rise Of MMP

The Experts



Kyle Bennett
[H]ard Talk
page 40



Alex "Sharky" Ross
The Shark Tank
page 39



Alex St. John
The Saint
page 14



Rob "CmdrTaco"
Malda
**The Department
Of Stuff**
page 90



Anand Lal Shimpi
Anand's Corner
page 38



Pete Loshin
Open Sauce
page 78



Joan Wood
Forward Slash
page 91



Lisa Lopuck
Site Scene
page 88



Chris Pirillo
Dialogue Box
page 77

Copyright 2002 by Sandhills Publishing Company. *Computer Power User* is a trademark of Sandhills Publishing Company. All rights reserved. Reproduction of material appearing in *Computer Power User* is strictly prohibited without written permission. Printed in the U.S.A. GST # 123482788RT0001 (ISSN 1536-7568) *CPU Computer Power User* is published monthly for \$29 per year by Sandhills Publishing Company, 131 West Grand Drive, P.O. Box 82667, Lincoln, NE 68501. Subscriber Services: (800) 424-7900. Periodicals postage pending at Lincoln, NE. POSTMASTER: Send address changes to *Computer Power User*, P.O. Box 82667, Lincoln, NE 68501.

Did you find the hidden CPU logo on our cover?
Turn the page for the answer.



Heavy Gear

- 16 **Extreme Hardware**
These Gizmos Don't Sing It, They Bring It
- 20 **The Big Showdown**
ATI & NVIDIA-Powered Cards Face Off
- 28 **ClearPC**
Cooler Master ATC-210 VX2
- 29 **Intel Pentium 4 2.4GHz**
AMD Athlon XP 2100+
- 30 **NVIDIA GeForce4 Ti 4200 vs.**
ATI Radeon 8500LE On The Cheap
- 31 **HP hn200w Wireless Gateway Router**
SyNet USB Pack
- 32 **Samsung SpinPoint SP8004H 80GB**
IBM Deskstar 120GXP 120GB
- 34 **Digiset DUO-64**
Lite-On LTR-40125S CD-RW 40X/12X/48X
- 35 **Adaptec XHub 4 Plus USB 2.0 Hub**
Sony CLIE PEG-T415
- 36 **Logitech TypeAway**
Digital Dream l'espion
- 37 **Apple iMac with SuperDrive**
- 38 **Anand's Corner**
Chipsets & Memory Standards
- 39 **The Shark Tank**
Sharpening My Teeth
- 40 **[H]ard Talk**
The Quickening



Hard Hat Area

- 42 **The PC Challenge:** Best Audio System
For Less Than \$1,500
- 44 **Swappin' Parts**
There's A Transformation Taking Place
- 46 **X-ray Vision:** 256-bit Encryption
- 48 **White Paper:** 3GIO
The Next
Generation Of Bus
Architecture



Loading Zone

- 70 **The Bleeding Edge Of Software**
Inside The World Of Betas
- 73 **Virtual Competition**
Connectix Virtual PC 4.3 vs.
VMware Workstation 3.0



- 75 **A So-So Extractor**
& **Typewriter Upper**
Easy CD-DA Extractor 4.6.9/Fanix
Software As-U-Type 1.2
- 76 **Eds For Your Webs**
Microsoft FrontPage 2002/Macromedia
Dreamweaver 4/Adobe GoLive 6.0
- 77 **Dialogue Box**
Snow White & The Seven Formats (+ 1)
- 78 **Open Sauce**
The Nature Of Computer Security: An
Interdisciplinary Approach

Caught In The Web

- 80 **Fringe**
- 83 **Google**
Searching The World, Finding Success
- 86 **Coder's Corner: XML**
Defining XML Dialects
- 88 **Site Scene**
Web Design Workflow: Part II
- 90 **The Department Of Stuff**
netradio.txt

- 91 **Forward Slash**
Hooked On Photonics

Digital Living

- 92 **Road Warrior**
Spring Blooms New Palms, BlackBerry
Gets A Voice, A New (Legal) Use For The
iPod & More From The Mobile Front
- 94 **At Your Leisure**
Plug In, Sit Back & Fire Away



Tips & Tricks

- 98 **Control Access**
Whip Your Work Into Shape
- 102 **Killer Hardware Tips**
Too Cool To Follow Rules

What's Cooking

- 104 **Technically Speaking**
An Interview With Greg Bildson,
CTO & COO Of Lime Wire
- 106 **Under Development**
A Peek At What's Brewing In
The Laboratory



Back Door

- 108 Q&A With Federico Faggin

Infinite Loops

Strange stats and other
oddball items from
computing's periphery.

15, 71, 80, 97, 100, 103



GREETINGS FROM SAMITLAND

I was just listening to the old George Gershwin tune "Summertime" on my PC Challenge machine (pg. 42): Summertime, and the livin' is easy / Fish are jumpin' and the cotton is high / Oh! Your Daddy's rich and your Ma is good lookin' / So, hush little baby, don't you cry . . .

And for whatever reason, it made me think of this month's ultra-cool Spotlight section. Sure, my daddy made a decent living as a doctor and my ma is good lookin', but sadly no amount of crying ever garnered me the BMX bike and the computers I desired so badly. I spent the summers of my youth yearning to become the best technofreak I could be, but I was never able to fulfill that dream at the time. Granted, it wasn't like I could just stroll down the road to pick up a new computer in Jos, Nigeria (I lived in West Africa from 1977 to 1987), so it would be unfair to heap the blame on my parents. Nor was I lazy, either. After all, I did manual labour most summers for up to two naira an hour (today that would be up to about 2 cents) through middle and high school. Nevertheless, I craved computers in the summer. To this day, summers make me think about upgrading my rigs. But first some nostalgia . . .

I fell in love with this topsy-turvy world of computers and technology nearly two decades ago. I fondly recall messing around on a (Timex) Sinclair ZX81 with the external 16KB memory module that made the computer lean back. The Commodore Vic 20, Commodore 64, Sinclair Spectrum, and TRS-80 were a few of the cool machines of the day that various friends had; we used to sit 2 feet away from 27-inch televisions with our home cassette decks (for saving and loading data) programming and playing games. I ended up buying an Oric Atmos (old French computer). It was awesome for learning BASIC, but ultimately failed against the C64 and Spectrum. Then just as MSX computers, the Commodore Amiga, and Atari ST were making a splash, it was time for this 17-year-old to hit America for university.

<Fast-forward a few years.> I got into the PC side of things more than a decade ago, and my upgrade cycles increased with consistent regularity. These days I usually do one major rig upgrade every summer. This desire harkens back to my youth—when I felt "the hunger." As the summer approaches, I find myself spending time digging deep into what's cool and what's coming before committing my hard-earned plastic to vendors all over the country. However, the Spotlight section this month will save me from the rigours of that research.

This month we take a very good look at next-generation chips and chipsets. Pat Gelsinger (Intel's CTO) mentions on page 58 that he predicts 30GHz CPUs by 2010. Did you get that? 30GHz. That's a long way from the computers that first got my attention all those years ago. I don't know about you, but I started daydreaming about Quake VII and intelligent voice recognition (to rescue my tired eyes and wrists), amongst many other things. Maybe by then computers will become man's new best friend, as my fiancée Julie coyly suggested to me the other day (she had braced herself for my gales of laughter, which she naturally received in waves). I digress.

Check out our Spotlight section this month to get the skinny on where we expect PCs to be in 2003 and what's coming down the pipe in the areas you're most concerned about. And that's not all. We have your usual heapings of stuff you love.

Well, that wraps up another month; keep that feedback coming. You'll soon have a way to be more interactive and discuss your comments with us more directly, but more on that later. For you "Star Wars" fans out there, don't miss our game coverage on page 94. I'll catch you here next month, same place, same time.

Samit G. Choudhuri

Samit G. Choudhuri, Publication Editor, CPU



Customer Service
customer-service@cpumag.com
Computer Power User
P.O. Box 82667
Lincoln, NE 68501-2667

Hours
Mon. - Fri.: 8 a.m. to 8 p.m. (CST)
Sat.: 8 a.m. to 4 p.m. (CST)

To Place An Order Or Change An Address

(800) 733-3809
FAX: (402) 479-2193

Subscription Renewals

(800) 424-7900
FAX: (402) 479-2193
<http://www.cpumag.com>

Authorization For Reprints

(800) 334-7458
FAX: (402) 479-2104

Editorial Staff

editor@cpumag.com
FAX: (402) 479-2104
131 W. Grand Drive
Lincoln, NE 68521

Advertising Staff

(800) 848-1478
120 W. Harvest Dr.
Lincoln, NE 68521
FAX: (402) 479-2193

The Gang

Editorial Staff: Ronald D. Kobler / Samit Gupta Choudhuri / Rod Scher / Christopher Trumble / Corey Russman / Kimberly Fitzke / Katie Dolan / Jennie Schlueter / Blaine Flamig / Raejean Brooks / Terry Dugan / Rebecca Christensen / Christine Mueri / Tara Weber / Sally Marek / Michael Sweet / Marty Sems / Chad Denton / Nathan Chandler / Calvin Clinchard / Kylee Dickey / Jennifer Seeman / Linda Rains / Dana Montey
Customer Service: Travis Brock / Alisha Lamb / Gretchen Elliott
Subscription Renewals: Liz Kohout / Connie Beatty / Matt Bolling / Patrick Kean / Charmaine Vondra / Miden Ebert / Kathy DeCoito / Stephanie Contreras / Theresa Sissel / Nicole Buckendahl
Art & Design: Lesa Call / Fred Schneider / Carrie Benes / Ginger Riley / Stacy Peters / Dave Fiala / Lori Garriss / Jason Codr / Andria Schultz / Erin Rodriguez / Ned Day / Lindsay Anker
Circulation Development: Kelly Richardson / Mark Anderson / Marcy Gunn
Newsstand: Mark Peery / Garth Lienemann / Dorene Johnson / Erik Egeberg / Christopher McGreer
Advertising Sales: Grant Ossenkop



April 2002 Corrections

"Leaders Of The Pack" (pg. 48): The \$3,699 price for the Alienware system did not include the price of the 22-inch NEC FE1250+ monitor. No wonder it was such a great deal. The full price including the monitor should have been \$4,491.

"Warm Up To Penguins" (pg. 99): The next release version of KDE was reported as being 2.3. The next release version will actually be 3.0.



Gotcha.
Here it is.

DVD changed the way you see movies.



Now it changes the way you make them.

Hitachi introduces the world's first and only multi-format DVD camcorder.

The incomparable Hitachi DVD camcorder brings the sophistication and simplicity of DVD to home movies. It records in two formats, DVD-R or DVD-RAM, and plays back on most computer DVD drives and most standard DVD players. It shoots still photos, too. You can edit on your computer and make copies of discs for friends with optional editing software. Or edit on the camcorder itself; DVD's simple navigation makes it easy to access recorded material. The Hitachi DVD camcorder. Home movie making will never be the same again.

Arriving in stores late spring.

HITACHI
Inspire the Next

ULTRAVISION
DIGITAL

UltraVision Digital is a registered trademark
of Hitachi Home Electronics (America), Inc.

www.hitachi.com/tv



Emmy® Award recipient for
pioneering camcorder technology

© 2002 Hitachi, Ltd.

In Software . . .



Tokyo's Avex Network distributes Avex-label CDs, among the latest to employ large-scale digital anticopy protection.

Titans vs. Pirates

In the clash between huge music corporations and audio pushers with apparent disregard for artists' intellectual property, most of us are caught in the middle. If you buy a CD, you expect that you'll be able to play it in standard CD and DVD players and in computer CD and DVD drives, right? And you probably believe you have permission, granted by the 1992 AHRA (Audio Home Recording Act), to make one copy of the CD for personal use. Sounds fair enough, but things are no longer that simple.

Those who have been following along know about the controversy surrounding Universal Music Group's decision to add anticopy technology to its CD soundtrack for the film "The Fast and the Furious." Universal used Midbar Tech's CDS (Cactus Data Shield) copy protection technology for the soundtrack CD, and now Japanese music distributor Avex Network is releasing 1 million CDs by three Asian pop stars using CDS technology.

Avex is using Midbar Tech's CDS-100 product to protect selected tracks on all three CDs. CDS-100 permits users to play protected tracks on standard CD players but not on PCs. By comparison, Midbar Tech's CDS-200 permits CD player and PC playback but prevents duplication. The company's CDS-300 product is due out sometime this year and should allow ripping to a PC or MP3 player while still protecting direct access to the original CD's contents. ▲

An Extremely Personal Firewall

Aimed primarily at technology companies with a sizeable mobile workforce, Check Point Software Technologies' new VPN-1 SecureClient software provides a personal firewall for a handful of popular PDAs and handheld PCs. As part of an all-star team that includes Compaq, Hewlett-Packard, Intel, and Microsoft, Check Point developed VPN-1 SecureClient to curb IT management stress levels.

The software enables employees to access company data via an 802.11-compliant, wireless Internet, dial-up, or Ethernet connection and provides a mechanism for centralized administration, monitoring, and streamlined software installation. VPN-1 SecureClient supports IP compression and works with HP Jornada 560 series Pocket PCs, HP Jornada 720 series handheld PCs, and the full line of Compaq iPAQ PDAs. ▲



Linux & Windows, Sittin' In A Tree

If you believe the hype coming from Lindows.com, the company's Linux-based LindowsOS software is poised to bring choice to PC users by letting Linux users run Windows-compatible apps on their machines.

Lindows.com was founded by MP3.com founder Michael Robertson and launched in 2001.

Calling LindowsOS a Windows "copy cat," Microsoft brought suit against Lindows.com. But a U.S. District Court judge for the Western District of Washington recently ruled that naming software to rhyme with Windows is not a crime. Representatives for Lindows.com successfully argued that the word

"windows" is a "generic term" and is therefore "not protected," so Microsoft can't claim the term as its own. To see the ruling, check out www.net2.com/lindows/ruling.pdf.

All of this makes for good press for Lindows.com, but it has yet to be seen whether LindowsOS is going to be a worthwhile addition to die-hard Linux users' desktops. The folks at Lindows.com were unwilling to divulge too much information about LindowsOS prior to the release of version 1, especially because the company released a LindowsOS sneak preview earlier this year, which didn't meet with unbridled enthusiasm. For one thing, LindowsOS

users always log on as root users, which most Linux and Unix users consider a no-no unless you're performing nitty-gritty system maintenance or installing new apps.

Using pieces of Wine, which lets Linux and Unix users run Win3.x and Win32 APIs on their machines, Lindows.com is developing LindowsOS so it will run any Windows application in the same way that it would run on a Windows 98/NT/XP machine. We asked Lindows.com how post-sneak preview development was coming along, but a spokesperson told us only that "Lindows.com expects V1 to be issued in the first few months of 2002." ▲

Putting The Pieces Together

Ever wish your shredding machine had an "undo" button? It probably never will, but ChurchStreet Technology Consultants (www.churchstreet-technology.com) is working on software that is bound to make certain accounting firm employees quiver in their wing-tipped shoes. The software set consists of three programs: Strip ID, Matching software, and Reconstruction software. Strip ID looks at scanned strips of shredded paper and reorganizes the fractured bits of text and images into manageable collections that the end user can reassemble.

According to ChurchStreet President and CEO

Cody Ford, development is proceeding at a good pace. "Preparing the strips for scanning is still the most labor intensive task," Ford explains. "Currently we scan the strips by attaching them to specially designed . . . reference sheets. . . . We [then] scan the strips on a flatbed to ensure nothing is curled up or misfed."

Once scanned, Strip ID assigns an electronic ID to each shred based on characteristics such as "length, color, white space, black space, font size, graphics, and indentations, to name a few," says Ford. After each shred has an electronic ID, the Matching software searches for possible matches and displays them for the user, who uses the

Reconstruction software to "drag and drop certain strips together, just as one would when completing a puzzle."

The hands-on reassembly process is necessary because, for example, law enforcement investigators don't want the software to automatically construct documents based on incomplete information. "We don't want to give the impression that we can fill in the blanks," says Ford, who adds that history and location are important factors in the reassembly process. ChurchStreet presently uses the software onsite and is perfecting its software for use in crime scene investigations and by the medical industry. ▲



Microsoft Under Pressure

If the company's song remains the same, Microsoft should be ready to release a beta version of Windows XP Service Pack 1 by the time you read this. WinXP SP1 should include support for Microsoft's Mira (handheld wireless devices that are slightly larger than a Pocket PC and have a relatively big screen) and Tablet PC, both of which are currently in production. Microsoft hasn't set a release date for Mira, but the final release version of WinXP SP1 and the Tablet PC should both be available by the end of summer. ▲

When Spying Eyes Are Smiling


Ever since the FBI announced it was using Carnivore software to monitor data packets as they pass through ISP servers (with court orders in hand, of course), the public has gained a heightened awareness about spying in the Internet age. Those who follow the spyware game know the technology behind e-mail and keystroke surveillance is nothing new. WinWhatWhere (www.winwhatwhere.com), for example, got its start in 1991 and moved from developing debugging software to releasing its titular product in 1993.

Multiple versions of WinWhatWhere's Investigator software have thrived in the workplace and generated scads of controversy mainly concerning the ethics of employee keystroke monitoring. While the war between corporate boss and civil libertarian rages on, WinWhatWhere has released Investigator 4.0.

Competitively priced at \$100 per seat and able to run on any Windows platform,

Investigator 4.0 adds a few new features to the mix, including Web cam monitoring (picture a Web cam peering at you all day long taking occasional snapshots that document when you're at your desk), key phrase alerts (whereby if you type a certain phrase for which the software has been programmed to look, the system monitor will receive an instant alert), and remote screen capturing.

There's also a feature WinWhatWhere dubbed "Hide and Seek," which makes it more difficult for users to find evidence that the software is running on their machine. So get back to work and be sure to participate in Hawaiian shirt day lest you receive a disapproving memo. As soon as we find software that lets you keep track of what the boss is doing, we'll let you know. ▲



The Investigator 4.0 Screen Capture feature lets managers keep a running log of intermittent workstation screen captures.

In Hardware . . .

10GB iPod Has Arrived

Apple's second-generation iPod doesn't look much different from the original, but with its 10GB capacity, the new model packs twice as many tunes. The new iPod also comes with a handful of new features, including a set of 20 EQ presets that let you play different music types (say, R&B, rock-and-roll, or classical) with equalizer settings optimized for the each type. Other new

iPod features aren't musically inclined at all, such as its ability to transfer a maximum of 1,000 names and addresses from contact information contained in an Entourage or Palm device or the Mac OS X Address Book. However, you have to download some free software to use the EQ and contacts features.

At 6.5 ounces, the new iPod weighs about the same as the old version, and it still uses IEEE 1394 (FireWire) for ultra-fast ripping. The 10GB model costs \$499, a hundred bucks more than the 5GB model, and is available online and in those friendly

neighborhood Apple stores that are beginning to crop up across the country. Windows 98SE/Me/2000/XP users can use Mediafour's recently released XPlay 5 software (www.mediafour.com/products/xplay) to use the iPod sans Mac. ▲



COURTESY OF APPLE

Atheros Chipset Breaks New Ground

Atheros Communications recently unveiled its AR5001X chipset, which is the first chipset to support all three IEEE wireless standards: 802.11a,

802.11b, and 802.11g. AR5001X includes three chips: AR5211 Multiprotocol MAC/Baseband Processor, AR5111 5GHz ROC (Radio-On-A-Chip), and AR2111 2.4GHz ROC. The chipset supports up to 16 54Mbps connections in the 2.4GHz to 5GHz frequencies, an improvement over single- or double-standard WLANs; for example, by itself 802.11b permits a maximum of three 11Mbps connections before performance starts to degrade. The chipset also includes hardware acceleration for improved security. When you read this, Atheros will be shipping the AR5001X in the United States. ▲



Let's Get Small

Scores of companies, research firms, and governmental agencies are currently developing applications using MEMS (microelectromechanical systems)-based technology, which integrates electronic and mechanical elements, sensors, and actuators on

silicon substrate during the fabrication process. (Take a look at page 52 of the March 2002 *CPU* for detailed coverage of the role MEMS plays in nanotechnology development.)

A seemingly infinite number of results are possible, including the 50-link

silicon micro chain that Sandia National Laboratories (www.sandia.gov) designed for the Department of Energy. "The chain was a resounding success," says micro chain inventor George E. Vernon. "One possible application for Sandia use



One benefit of MEMS-based technology is a display Microvision developed that lets you see the display, as well as what lies beyond.

The Telltale LED

Edgar Allen Poe wrote about the telltale heart and its imagined beat that betrayed a killer's dark secret, but gathering information by analyzing certain LEDs (light-emitting diodes) from afar is apparently real. So say Joe Loughry of Lockheed Martin Space Systems and David A. Umphress of Auburn University in their paper, "Information Leakage from Optical Emanations."

According to Loughry and Umphress, the risk that a blinking LED could expose an actual data stream is low for Class I LEDs, which indicate

whether a device's power is on; and risk is medium for



Class II LEDs, which indicate a device's activity level. Class III LEDs, however, such as those used in modems, can

reveal the actual data as it is transmitted.

The authors write, "If the correlation [of LED signals

with corresponding data] is sufficiently good, then from analysis of Class III optical emanations it is possible to recover

the original data stream."

Loughry and Umphress tested a "wide variety of devices" and found "exploitable compromising emanations in 36% of devices tested." In short, someone in a building across the street could potentially tap into a vulnerable data stream with a modicum of effort.

To read the entire paper, visit Loughry's Web site (applied-math.org) and select a PDF or compressed version. Official publication is scheduled for the August issue (Vol. 5 No. 3) of the ACM's (Association for Computing Machinery) journal, *ACM Transactions On Information And System Security*. ▲

is in locking devices. . . . My research on macro scale chains indicates the smallest conventionally machined chains are used in cameras to operate the shutter. This is certainly one possible industrial application [for the micro chain]. Almost any application which requires multiple or serial parallel events to occur could use this chain."

Microvision is also on the cutting edge of MEMS technology, having recently released a MEMS-based, full-color display with SVGA (800 x 600) resolution that measures just 1.5 inches across. The company's Nomad Personal Display System is the first

commercially available MEMS-based display, and it provides an "augmented reality" view in which you can see behind the image that's displayed. (See page 88 of the May 2002 *CPU* for details about Nomad.)

According to Microvision CTO Casey Tegreene, small, full-color, SVGA displays will make it possible to replace complete paper manuals with electronic versions for various industries where portability is a significant issue. Pilots and medical personnel, for example, could benefit from having such manuals readily available, perhaps even on a flip-down personal

viewer. Tegreene believes MEMS-based displays will find their way into all sorts of mobile devices and says Microvision is currently working on ways to make them smaller and less power-hungry.

Another application that holds a lot of potential for MEMS technology, says Tegreene, is 2-D coding. You can find examples of 2-D coding, which looks like a three-ring target surrounded by a bunch of dots, on packages processed by any major carrier. With names such as Maxicode, QR Code, and Snowflake, 2-D coding is an alternative to 1-D coding, such as bar codes,



Sandia National Labs used MEMS technology to create this micro chain, which measures 50 microns between each link (a human hair is approximately 70 microns in diameter).

which hold relatively little data, and is less expensive than competing technology, such as RF (radio frequency) coding. According to Tegreene, programmers can embed JPEG images and security code within a 2-D

code, thereby enlarging the range of possible applications to include security badges. For more information visit the Microvision site (www.microvision.com) or the MEMS Clearinghouse (www.memsnet.org). ▲


Internet . . .

Break On Through

Forget View Source. As long as you stay within the legal lines, you can use BCF Technology's new Site Raider software to capture code and, as a Site Raider spokesperson explained to us, "deconstruct ANY Web site." The software lets you view all of a site's code, including hidden or embedded code, and even lets you perform a security risk analysis.

With Site Raider you can point at a Web site's content and run the clip command to obtain a copy of the code. This includes the code behind dynamic functions such as user logins, financial transactions, and online forms. If you click Generate Source Code, Site Raider generates bug-free code in the language of your choice, including Java, JavaScript, ASP, and Cold Fusion.

Site Raider runs on Windows Me/NT/2000/XP, and a free 10-day trial is available. To truly see everything there is to see, you'll need to shell out \$429 for Site Raider Professional. BCF Technology also sells Site Raider Standard for \$299 and Site Raider Light, which offers more functionality than we expected, for \$149. ▲



Is Site Raider legal? Well, that all depends on how you use it, as the disclaimer on the Site Raider site states.

Search Engines Tilt At Google's Windmill

Although SearchOnline.info (www.searchonline.info) doesn't reinvent any search engine paradigms, it does search all major indices (not just one, like some of the big sites do) and it puts an ingenious spin on presenting search results. Results are prioritized in the same fashion that other search engines provide, but the site lets you manipulate

the results further. For example, you can opt to group duplicate URLs or sort by popularity, contents, domains, URLs, or titles. And if Search-Online delivers more than one hit for the same Web site, only the site home page is listed and a plus sign appears beside the URL. By clicking the plus sign, you can see all of the relevant hits within that Web site with-


out altering the currently displayed list.

We tried a few identical searches on SearchOnline and Google, and although some Search-Online results were off-topic duds, most were equal in

value to what came up on Google. And, just like Search-Online claims, we found more useful and various news hits among the SearchOnline results. We didn't like Search-Online's single-box approach to searching (as opposed to Google's multibox advanced search feature) but did appreciate SearchOnline's greater number of Boolean search options and the fact that multiple same-site results didn't each count as a separate hit.

Search results also include ordinal number hit rankings for other sites. For example, "Yahoo! 1st" identifies a result that Yahoo! lists first for the same search. SearchOnline does this for AllTheWeb, AltaVista, DirectHit, LookSmart, Lycos, MSN, NorthernLight, Teoma, WiseNut, and Yahoo!.

Another search engine, Teoma (www.teoma.com), debuted shortly before press time. Like SearchOnline, Teoma is poised to knock Google down a rung with its subject-specific search engine. Rather than providing results that are simply popular or plainly relevant, Teoma's forte is delivering results that same-subject pages reference. We tried the same searches on Teoma that we tried on SearchOnline and Google; not only did Teoma deliver far fewer duds, but every one of its top 19 hits qualified as what we'd consider highly relevant. Considering that with every search Teoma also serves up excellent search refinement tips and a handful of mostly relevant resource links, Google might be facing some stiff competition. ▲



SearchOnline.info offers a new twist on searching the Web.

Spam Never Sleeps

No matter how advanced our antispy technology becomes, the stuff just keeps popping up uninvited like dinnertime telemarketer calls and pseudo-high-tech gadget catalogs. Of course, some folks eventually get fed up enough to stand up, lean out of their windows, and shout, "We're not going to take it anymore!" But spam (the e-mail version, not to be confused with the Hormel meat product), is a slippery thing, and some courts have begun hearing some interesting cases.

Take, for example, the California Supreme Court, which recently decided to review a lower court ruling that companies have the right to sue folks who send spam or other unwanted e-mail to employees via the company's e-mail system. Intel brought the suit against ex-employee Kourosh Kenneth

Hamidi, who sent six e-mail messages over two years to all Intel employees except those who e-mailed back requesting that Hamidi send them no further messages.

Apparently Intel high-ups didn't appreciate the messages' negative, anti-Intel tone. Litigators for Intel unearthed the arcane 18th century "trespass to chattels" legal doctrine, claiming that trespassing in order to interfere with a company's property is applicable to spammers sending e-mail through the company server. The EFF (Electronic Frontier Foundation) calls it "unwise" to apply this interpretation to spam, and the ACLU (American Civil Liberties Union) also disapproves.

Opening briefs should be delivered by the time this issue hits the newsstands, so jump online to see how it turns out. The ACLU of Northern

California is following the case, and you can read the entire Intel v. Hamidi brief at www.aclunc.org/expression/intel-brief.html.

Lest true spammers take the EFF and ACLU interpretations as tacit approval for further spending, they should pay attention to other news. For example, the FTC (Federal Trade Commission) recently announced numerous initiatives: one targets spam and Internet fraud in northwest United States and southwest Canada; another targets spam that appeals to post Sept. 11 patriotism but delivers worthless products. The FTC has also been cracking down hard on Internet chain letter and pyramid scams. To help out, forward unwanted spam to uce@ftc.gov and add to the database of more than 8.5 million messages. Now that's a lot of spam. ▲

New On The 'Net

Looking for some new surfing destinations? Here's a sampler of the many sites that recently hit the Web.

Bargain Basement PDAs

Think Pricegrabber for PDAs, and you'll get an inkling of what bargainPDA (www.bargainpda.com) has in store. The difference is that bargainPDA is dedicated to all things to do with, well, bargain PDAs, not just PDA prices. On the site, you can click Price Watch to search for a product type, manufacturer, and/or model. The search engine then delivers up-to-date product specs, pricing, and reviews. Better yet, click Coupons & Rebates to access a list of links directly to sites offering deals; there are tons of deals there from distributors such as BestBuy.com, ComputerGeeks.com, and Handspring.com. Or click Resources to find links to manufacturer and retailer Web sites and PDA user groups. Overall, the site offers a nice selection of software, used PDAs, and reviews.

Complain, Complain

Has some company, online or otherwise, done you wrong? Would you like to lodge a complaint but you don't have the time or resources to do the job justice? For \$5.95, Complain Complain.net (complaincomplain.net) will send a letter to the dastardly company's owner, CEO, president, or other high official. Complain Complain guarantees (in capital letters, no less) that you will, at minimum, receive an apology letter from the company. Complain Complain started up in Detroit, Mich., in March 2002 and, by launching this site, has since grown wider in scope. ▲

IOS Upgrades Available Online

Before you send another motherboard to the landfill, consider upgrading the BIOS and giving your PC a new outlook on life. Here are a few recently released upgrades. Check out www.smartcomputing.com/cpumag/jun02/bios to see the entire upgrade list.

Manufacturer	File/Date Available	URL
ABIT	LOAD105 version 1.05 C4 (03/29/02)	www.abit.com.tw/abitweb/webjsp/english/download.jsp
ABIT	NV784, HPT 37x BIOS version 2.0.1024 (03/07/02)	www.abit.com.tw/abitweb/webjsp/english/download.jsp
ABIT	ST6E8B, HPT 370 RAID BIOS version 2.0.1024 (03/05/02)	www.abit.com.tw/abitweb/webjsp/english/download.jsp
AOpen	AX4BS-V version R1.05 (03/11/02)	www.aopen.com/tech/download/mbbios/ax4bsv.htm
AOpen	AX4BS Pro version R1.05 (03/08/02)	www.aopen.com/tech/download/mbbios/ax4bspro.htm
AOpen	Fortress 1100 BIOS version R1.02 (03/06/02)	www.aopen.com/tech/download/mbbios/ft1100.htm
AOpen	MX36LE-UN version R1.07 (03/12/02)	www.aopen.com/tech/download/mbbios/mx36leun.htm
AOpen	MX36LE-U version R1.07 (03/12/02)	www.aopen.com/tech/download/mbbios/mx36leu.htm
AOpen	MX34-U version R1.05 (03/11/02)	www.aopen.com/tech/download/mbbios/mx34u.htm
AOpen	AX4B Pro version R1.03 (03/08/02)	www.aopen.com/tech/download/mbbios/ax4bpro.htm

Compiled by Steve Smith

You Get Paid For That?

Those of you with Y chromosomes will think you died and went to geek heaven if you land this choice IT job: director of Internet technology for Playboy.com. Working from the Chicago headquarters (that's right, the main bunny hutch), you would work with the business and editorial staffs to create new interactive features and manage out-sourced services.

Sure, you'll need to know Java, Flash, HTML, and C++, as well as be able to identify reliable vendors, but we're pretty sure the perks make this job worth it. After all, someday a bunny may get lost in the Playboy tower and meander into the IT department by mistake. ▲


WiFi Hot Spots Start To Simmer

Waiting for the day when you can whip out your notebook in a coffee shop or airport and log into the nearby wireless LAN, grab e-mail with your latte, or check forecasts at your destination while you wait for the flight? Anyone who has an 802.11b (WiFi) LAN card in his portable PC

dreams of the day when he can log into wide zones or "hot spots" of wireless signals directed into public areas. This may be the year it finally happens.

According to wireless provider MobileStar, the number of WiFi hot spots will increase dramatically in 2002, with retail stores (mainly wired coffee shops) leading the way,

followed by hotels and airports. Expect more than 5,000 hot spots in operation worldwide. Hotels will be a major growth area for this market, as will the airline courtesy lounges for regular business travelers. The Yankee Group is predicting there will be 15,000 public places from which to pull the Web out of the ether by 2005.



Certifiably Techy

Just when you thought you had seen your last exam, it turns out that many IT professionals are lining up to take more classes. As the job market tightens and is flooded with potential employees, companies are looking harder at skills certification when hiring.

It may not be enough that you have 10 years of experience making a corporate tech

infrastructure run smoothly; your next boss may want a piece of paper that says you know your way around a Cisco router or how to implement security measures. According to Foote Partners LLC, skills certification increases an IT worker's pay more than 8% in raises and bonuses. So it makes sense that people are flocking to these courses, each of which

can set you back \$1,200 to \$11,000 per certification.

The certification business is projected to soar from \$13 million in revenues in 2001 to \$20 million in 2005, according to IDC. With so many jobs and so much money at stake, it isn't surprising that tales of piracy and incompetence plague the field. The industry is trying to address

these problems with a consortium of tech companies ominously named the ITCSC (Information Technology Certification Security Council; www.certsecurity.org). It will confer seals of approval on testing companies to reassure employers that these costly but ubiquitous certificates are worth the paper on which they are printed. ▲

You just stepped off the plane, yet
you can't wait to sit down again.



The Lincoln LS at Hertz.
The luxury car you're yearning for is now part of
the new Hertz Prestige Collection.

Rent the Lincoln LS and sit in the lap of one of the
luxurious performance vehicles included in the new Hertz
Prestige Collection. Simply reserve the one you want
(available at select U.S. and Canadian locations) for your
next business trip or vacation. For reservations, call your
travel agent or the Hertz Prestige Collection reservation
line at 1-800-654-2250, or visit us at hertz.com.
And you'll get special services that are exclusive
to our Prestige Collection. Reserve your seat today.

hertz.com

Not all vehicles, vehicle equipment
and services available at all locations.

® REG. U.S. PAT. OFF. © HERTZ SYSTEM, INC., 2001/09/01/01

Hertz
Prestige Collection

by Alex St. John

The Rise Of MMP



Alex St. John was one of the founding creators of Microsoft's DirectX technology. He is the subject of the book "Renegades Of The Empire" about the creation of DirectX and Chromeffects, an early effort by Microsoft to create a multimedia browser. Today Alex is president and CEO of WildTangent Inc., a technology company devoted to delivering CD-ROM quality entertainment content over the Web.

I recently found the original proposal for DirectX that I wrote back in 1994 containing this quote:

"Connectivity combined with other technologies may fundamentally change the business model of the game industry from writing \$50 throw away applications, to building vast extensible server-based game universes for many players."

I was the first game evangelist for Microsoft back then, and I believed that the addition of standard networking APIs to the Windows operating system would enable a mass market for multiplayer video games. Doom from id Software was the gaming craze back then, so I made a trip to Texas to meet John Carmack for the first time. John lent me the source code to Doom to try to port to Windows. The first attempt to port Doom to Windows was futile and even with a new graphics technology from Microsoft called WinG, Windows did not support the media capabilities required to play Doom. At the time, I did not have engineering resources of my own, but a Program Manager named Gabe Newell lent me an engineer to work on the project. Gabe Newell later retired from Microsoft to form Valve Software, where he created the game Half-Life based on id's Quake gaming engine. Doom went on to become the first DirectX game. We used it as the benchmark for the DirectX 1.0 feature set. Doom95 was one of the best selling Windows 95 applications when it was released. On Halloween 1995, Microsoft hosted the multiplayer Doom finals at a party called Judgment Day where Dennis Fong (aka. "Thresh") became the first video game playing star in the industry.

On that trip to Texas I also visited Origin, where I had dinner with Richard Garriot and several of his most talented engineers. We talked extensively about the possibility of creating a multiplayer version of Ultima 8. Richard was not convinced that it was a good idea at the time, but agreed to send three of his engineers to Microsoft for a month to work

on porting Origin games to the early DirectX technologies. These engineers, working closely with the DirectX development team, ensured that Origin's leading-edge games could indeed run on Windows using DirectX, in many cases faster than they could run under DOS. It was during this visit that these engineers also adapted the Ultima 8 engine to create "Multima," a 256-player experiment that went on to become Ultima Online.

DirectX 1.0 included a new API we had created called DirectPlay to enable game developers to easily create multiplayer games. DirectPlay 2.0, released in 1996, added lobby services and server architecture designed to enable developers to create Massive Multiplayer Games. DirectPlay was used heavily

for multiplayer games within Microsoft but was not adopted widely by the game development community. Today Microsoft purportedly has a large team working on developing DirectPlay 9.0.

Massive Multiplayer Games are all the rage in the game industry these days. Everybody is busy building a giant

MMP world to compete for the same 250,000 hardcore gamers. Many game publishers seem to believe that building a giant expensive MMP game is the only way to have a profitable online gaming strategy. When the first MMP games were created, game developers had no idea exactly what kinds of gameplay would be popular or what kinds of problems they would have with an MMP game community. (For example, 80% of support calls for some MMPs are complaints from players about other players "bothering" them.) The early MMP games like Ultima Online and Everquest are vast persistent worlds that let thousands of players occupy the same world at the same time, even though each player only interacts with a small subset of the active players. I am of the opinion that these early experiments in MMP game design will eventually give way to much lighter weight, less infrastructure-bound MMP game designs that benefit from the experience

**It's [eBay] the ultimate
Pokemon/Magic The
Gathering trading game:
competitive bidding and
trading are of course another
excuse to e-mail. eBay is one
of the Internet's great new
game companies.**

of the early pioneers to create game designs that capture the value of persistent worlds shared with multiple players, but resolve many of the game design challenges associated with these worlds more elegantly.

I also suspect that there are many other types of MMP game designs that have great potential to be popular and profitable on a very large scale but will require some adventurous developers to pioneer. When people tell me that MMP games are just for a small niche of hardcore players, I like to list my top five money-making MMP games that most people don't realize are genuine MMP games. Our business is so wrapped up in 3-D engines that other kinds of MMP games have been cropping up on the Internet without notice.

#5: eBay. It's the ultimate Pokemon/Magic The Gathering trading game: competitive bidding and trading are, of course, another excuse to e-mail. eBay is one of the Internet's great new game companies.

#4: E-mail/Chat. What do AOL's 33 million users spend their time doing online for \$19.95/month? E-mailing and chatting with each other in TEXT! Yes that's the primary use of an AOL connection. You think browsing is a relevant portion of that? Nope, they're just looking for funny or interesting content to forward to each other as conversation starters.

#3: Adult Content. This is a 2.2 billion dollar subscription business online. Streaming peep shows in which the audience directs a performer's actions via chat in a streaming video window offers the ultimate in 3-D realism and interactivity without investing three years to develop a leading-edge game engine.

#2: Online Gambling. Has anybody else noticed that the only banner ads on gaming sites these days are for online casinos? They're making billions. I have a friend who runs an offshore casino-hosting business. He processes 450 million dollars in transactions a month; this with 2D card games that everyone else gives away to "casual gamers" for free. It's illegal, you say? Skill-based games for cash prizes are not illegal, but nobody does it on a very large scale.

#1: E-Schwab. This online gambling game matches millions of players against the aggregate moves of every other player to maximize their personal portfolios. Players gamble real money and pay \$29.95 to the house for every move. E-Schwab manages billions of dollars in assets for its players. A remarkable achievement for a game with a text-based user interface!

These are some good examples of online entertainment content that has a mass audience and can make vast sums of money without looking like Quake or needing EverQuest infrastructure. People like to point to the failure of EA's Majestic as a reason to believe that other adventurous forays into MMPs are likely to fail, but there are an infinite number of wrong game ideas to be had and a small number of right ones. At least EA has had the courage and initiative to explore them. I'm betting EA has a winner in The Sims Online.

Within the game industry, there are also some interesting examples of games that should be considered MMP but aren't. Games like Quake Arena from id Software rely on their fan base to set up and manage their multiplayer infrastructure for them. Instead of selling subscriptions, they sell retail expansion packs. Services like


GameSpy provide chat and lobby services for their players outside the game. The collective result is an experience that captures much of the value of an MMP community without the infrastructure costs and headaches. Blizzard's Diablo II dynamically joins players for gaming sessions and keeps character profiles online, again capturing many of the valuable elements of MMP game play without the plumbing overhead. I'm personally very excited about the imminent arrival of Neverwinter Nights, a new game from the folks [at BioWare] who brought us Baldur's Gate, that will allow players to create unique game levels and invite their friends online to play in them.

Most major game publishers have turned their attention to creating console games for the next few years, but I believe that they may be surprised to discover that the console momentum may not be as strong this time as it has been in the past. Current generation consoles, even the Xbox, are online-challenged, and consumers may have had their fill of isolated 3-D glitz. The Internet has given them a strong appetite for community. ■

*Direct your pure thoughts to
TheSaint@cpumag.com.*

Infinite Loop

Voice Mail Maze



The world may never know how many licks it takes to get to the center of a Tootsie Pop, but we set out to learn how many buttons you have to push to get to a live person on a technical support line. We called five companies on an ordinary working Monday between 2 and 3 p.m., pretending to seek support for a desktop PC. We documented the number of clicks required to get from each company's home page to the appropriate phone number, the number of voice mail layers and options that stand between dialing the phone number and speaking with a live person, the time it took to get that far, and how long we waited on hold.

	Apple 800-275-2273	Compaq 800-652-6672	Dell 800-624-9897	Gateway 800-846-2301	HP 208-323-2551
Clicks required to find phone number on company Web site	1	2	1	1	3
Voice mail layers before getting put on hold	4	2	5	3	6
Voice mail options, combined total for all layers	3	6	3	2	2
Time elapsed before getting put on hold (mins:sec)	11	17	7	6	7*
Time spent on hold (mins:sec)	1:20	2:17	1:23	1:25	1:30
	0:05	0:30	0:15**	0:48	0:21

* Hewlett-Packard's voice mail system uses voice recognition.

** After 15 seconds, Dell's system required an item number, which our fictional PC lacked.

EXTREME HARDWARE

These Gizmos Don't Sing It, They Bring It



Timex Watches With SpeedPass

If you like to patronize places with signs reading, "Eat Here—Get Gas," have we got a wristwatch for you. But first, some background on what petroleum and McDonald's fries have in common. For a few years now, ExxonMobil has offered the SpeedPass system (www.speedpass.com). SpeedPass lets reportedly 5 million users buy gas by waving a key fob over a sensor at an Exxon or Mobil gas station nationwide, which charges the purchase to a credit card. Ronald McDonald liked the idea, so Chicago SpeedPass users also can use it at roughly 400 Mickey D locations. Other businesses and services could follow suit. Timex, meanwhile, is testing six styles of watches incorporating the technology (www.timex.com). They probably won't be available this year but may eventually retail at \$30 to \$50, says spokesman Jim Katz. If your SpeedPass watch gets ripped off, Katz says, you handle it like a credit card: Your financial liability is limited, and you can call a number to deactivate your account. These watches could only be cooler if SpeedPass charged Bill Gates' credit card instead of our own. Like he'd notice.

It's time to beat the heat once again, but this wouldn't be

Extreme Hardware if we didn't give you some odd ways to do

it. For instance, if your Hitachi notebook gets your lap a little

warmer than you'd like, it may be crying out for another glass

of water. And if your next skillet fears it's going to scorch your

grilled cheese and banana pepper masterpiece, it might tell your stove

to chill for a moment. Finally, if you'd rather cool off on the open road,

don't forget to pay for gas and a burger first with your wristwatch.

by Marty Sems

Pentax DIGIBINO DB100

It's a common dilemma: You're scanning the rest of the apartment complex with your binoculars for hours, eating Cheetos, and you finally see something you'd just die to have as a photo. Well, OK, maybe it's only a dilemma common to our pub editor, Samit (rimshot). We suggest Pentax's new DIGIBINO DB100. It's a mutant combination of a 7X CCD digital binocular and a 3.3X zoom single-lens digital camera. The camera part has a progressive CCD, a 1/8000-second shutter speed, 0.8 megapixel (1,024 x 768) resolution, and 16MB of built-in memory. No, you can't take stereoscopic or 3-D photos with it, but at least you can catch the action as it happens. Scope it out this June or July for about \$299 to \$349 (www.pentax.com).





Panasonic Body FX

Admit it, you want to feel Halo's explosions rumble your kidneys. And you know you want the feel of fingers caressing your shoulder blades as Lil' Kim raps something unprintable. You want Panasonic's Body FX. It's a massage pad that not only can stroke away your tension with inflatable "fingers" and its own separate air compressor, but also can pound your lower back with a big, bassy transducer speaker. The really stimulating part is that the Body FX's speaker will respond to any audio source you stick into its miniplug audio jack (headphone splitter included), such as your stereo, game console, or DVD player. You can buy a \$400 Body FX with a seat from Brookstone (www.brookstone.com) in June, or a \$350 seatless one in August or September. Panasonic says either version is usable on a chair, sofa, or bed. "We can think of a lot of uses for this," says spokesperson Dan Silver. Whatever your pleasure, make the earth move during your next aural experience with the Body FX (www.panasonic.com).




Hitachi Flora

Hitachi's about to water its Flora (www.hitachi.com). Breaststroking ahead of other major manufacturer factory options that typically dive first into the deep end of desktop systems, water-cooling appears poised to make a splash in a corporately sanctioned notebook. If all goes as planned, Japanese mobile users should be able to buy a watercooled Flora Pentium 4 notebook (for an undetermined price) in Q2 or Q3. Pumping liquid through a CPU cooling block and then through a radiating panel safely sealed behind the LCD, the Flora's cooling system should cool a 30W processor with an almost imperceptible 26dB to 30dB of noise. And Hitachi says the obvious weak link, the LCD hinge area, shouldn't wear out for about 20,000 cycles. In the high-tech ecosphere, that's as eternal as the ocean.

Smart Pan RF Smart Cooktop

Our readers like mods that let them monitor temperatures. Fair enough. Here's one for your skillet. If you're the type of gomer who assumes that one fries spaghetti, this development could keep your house from burning down. Digital Cookware's Smart Pan RF is a \$95-to-\$98 nonstick skillet with a programmable LCD temperature gauge in the handle. This is neat in itself, but it ain't soup yet. If the burner gets too hot for your pathetic attempt to fry a Hot Pocket, the pan sends a 433MHz RF signal to the control module installed in your 120v or 240v stove, which adjusts the burner to maintain the temp you want. And if your meal commits flambé or something and hits 450 degrees, the Smart Cooktop will shut the burner off. Digital Cookware says the Smart Cooktop control module may show up in name brand ranges for about \$10 extra (www.digitalcookwareinc.com).





See more of your ideas. Less of your monitor.

Introducing the world's first, most complete line of thin-frame options from the best-selling brand of flat panel monitors. More picture. Less frame. Nothing to distract you from amazing color and clarity, or take away from edge-to-edge flicker-free viewing. At just 18.5 mm per side, it's the industry's thinnest frame. Plus, the MultiSync® LCD1850X, LCD1550X, LCD1550M and LCD1550V offer a choice of sizes and colors; touchscreen and protective glass technologies; height-adjustable stands; portrait and landscape viewing; and an endless array of multimedia options. And their sleek designs provide virtually seamless viewing in multi-monitor applications.

So broaden your horizons. Choose Ambix™ for digital/analog dual input technology that ensures long-term compatibility. XtraView+™ for wide angle viewing. OmniColor™ technology for precise color values. Or LiquidView™ for high-resolution readability. It's innovation that lets you see more of what you want. Less of what you don't.

Learn more at www.necthinframe.com or call 888-NEC-MITS.

SEE MORE.™

NEC ranked as the #1 flat panel monitor branded vendor in the Stanford Resource—iSuppli Flat Panel Monitor Quarterly Report, Q2 '01. MultiSync is a registered trademark, and Ambix, XtraView+, OmniColor and LiquidView are trademarks of NEC-Mitsubishi Electronics Display of America, Inc.

©2002 NEC-Mitsubishi Electronics Display of America, Inc.
All rights reserved. Simulated images in monitors.



The Big Showdown

ATI & NVIDIA-Powered Cards Face Off

There have been several great rivalries between video card manufacturers over the years, including Diamond Multimedia vs. Hercules, 3dfx vs. NVIDIA, and now, NVIDIA vs. ATI. The battle is a bit like a boxing match. Power users anxiously watch each contender, waiting to see which one will overpower the other with brute strength, speed, and a touch of grace.

Overall, today's video cards are unbelievably powerful, but so were the cards that came out last year, and those the year before that, comparatively speaking. Remember when 32MB of video RAM seemed ostentatious? Now, cards with 128MB of RAM are the rage.

There have been other significant changes in the industry lately, as well. ATI finally decided to share its video chips with retail video card manufacturers. The new ATI-powered video cards are already available from such manufac-

turers as Hercules and Gigabyte.

NVIDIA has always worked with numerous manufacturers to get its GeForce GPU into users' hands. This year, however, NVIDIA released four new GeForce GPUs at the same time, unlike past years when NVIDIA has released a flagship GPU fol-

lowed by other versions during the course of the next few months. The powerful GeForce4 Ti 4600 is this year's top-of-the-line GPU, with the GeForce Ti 4400 being a lighter, slightly less expensive, version. These GPUs incorporate the kind of high-powered octane that make users lustfully reach for their wallets with one hand while wiping spittle from their cheeks with the other.

NVIDIA is smart enough to realize that not every user desires or means to casually drop \$400 on a video card, no matter how sloppily he's

drooling over himself. Thus, NVIDIA released the budget-priced GeForce4 MX 440 and the ultra-budgety GeForce4 MX 420. The "GeForce4" moniker might lead one to believe the MX 440 and MX 420 cards are a notch above GeForce3 video cards released last year, but this isn't the case. GeForce4 MX cards don't support all the DirectX bells and whistles that GeForce3 Ti cards do, and GeForce3 Ti video cards can outperform a GeForce4 MX video card.

We reviewed 2-D/3-D cards from ATI, Gigabyte, Hercules, Leadtek, and Visiontek. We also requested video cards from Asus, PNY, Gainward, and ABIT but didn't receive products in time for testing. We also reviewed a video card from Matrox and a workstation video card from 3DLabs to detail how they're different from the usual 2-D/3-D video cards.

How We Tested

We tested our video cards on a 1.8GHz Pentium 4 system with 512MB of RDRAM, running Windows XP Pro Edition. We ran three sets of tests on each card, including 3DMark2001, a Quake III time demo, and two Serious Sam time demos (Karnak and Metropolis). We ran the Quake III demo at three resolutions, noting the frame rate at each resolution. We ran the Serious Sam demos at 1,600 x 1,200 to see how each video card could handle the game's high-graphics workload. We tested the Matrox card using MadOnion's 2-D graphics benchmark, SYSmark2002. Each video card in our roundup has a maximum resolution of 2,048 x 1,536 and a maximum 32-bit color depth.

OK, enough small talk. We know you want to know which video card is the big shot and which ones should wither away on shopkeepers' shelves. Let's get down to business.

ATI Radeon 8500 128MB

You can find everything you really need to know about the ATI Radeon 8500 128MB in its product name. The card is an improved version of the original ATI Radeon 8500 video



Radeon 8500 128MB

\$299

ATI

(905) 882-2600

www.ati.com



card; it still has the Radeon 8500 GPU but ATI doubled the card's memory to 128MB of DDR-SDRAM. That should be more than enough RAM for even high-end gamers. The Radeon 8500 128MB's refresh rates are from 85Hz to 200Hz, and the card supports Win98/Me/2000/XP.

Overall, the 8500 128MB performed nearly as well in our tests as the Visiontek and Leadtek GeForce4 4600 video cards we reviewed. The card's 3DMark2001 score was 8,496, third best among the video cards in this review.

The card also handled the Quake III demo pretty well at every resolution. Its frame rate at 800 x 600 was 198.6fps, and at 1,024 x 768, the frame rate was 186fps. These are extremely fast frame rates, although they were still about 10fps to 15fps slower than the two GeForce4 Ti 4600 cards we tested. The Radeon 8500 128MB fell far behind the GeForce4 Ti 4600 cards when we increased the resolution to 1,600 x 1,200, however, charting a frame rate of 108.4fps compared to the two GeForce4 4600 cards' average of about 158fps.

The Radeon 8500 128MB handled the Serious Sam demos well enough, although it was significantly slower than the Leadtek and Visiontek video cards. During the Karnak demo, the Radeon 8500 128MB's frame rate was 66.3fps, and the frame rate during the Metropolis demo was 75.9fps. Although these frame rates were slower than the GeForce4 Ti 4600 video cards, the gameplay was still solid.

The Visiontek and Leadtek's video cards outmuscled the ATI Radeon 8500 128MB, but the 8500 128MB does have a price advantage over those cards, costing just less than \$300. You'll have to spend about \$400 for a GeForce4 Ti 4600 card.

Gigabyte AP64D-H

When you think of Gigabyte, you probably think of motherboards. That may be about to change now that Gigabyte has partnered with ATI to make ATI-powered video cards. Gigabyte's AP64D-H is one of the best cards in Gigabyte's lineup, with a Radeon 8500 GPU and 64MB of DDR-SDRAM. The card has refresh rates up to 200Hz and works with Win98/Me/2000/XP.

Overall, the AP64D-H proved to be a decent video card. It practically kept pace with

ATI's Radeon 8500 128MB card, even with half the memory. The AP64D-H's 3DMark2001 overall score of 8,169 isn't bad by any means, although three other video cards scored higher.

The AP64D-H rumbled through the Quake III demo, posting good frame rates at every resolution. Its frame rate at 800 x 600 was an impressive 197.4fps, although that was only the fourth fastest frame rate in our roundup. The AP64D-H's frame rate dropped only slightly when we increased the resolution to 1,024 x 768, maintaining a frame rate of 183.3fps. At 1,600 x 1,200, the card dropped to 103.8fps, but that's still more than fast enough to run Quake III smoothly.

The AP64D-H's frame rates in the Serious Sam demos were respectable, although the scores again finished behind the ATI Radeon 8500 and the two GeForce4 Ti 4600 cards we tested. The AP64D-H's frame rate in the Karnak demo was 63.9fps, while its frame rate in the Metropolis demo was 70.8fps.

The card managed to keep up with ATI's Radeon 128MB card for the most part, despite that card's 64MB memory advantage. Really, there's not enough of a performance difference between the Gigabyte and ATI cards to justify paying for the extra memory (at least not until games arrive that can utilize all that memory), so the Gigabyte AP64D-H is ultimately a better deal.

Hercules 3D Prophet FDX 8500 LE

The 3D Prophet FDX 8500 LE is less expensive than most new video cards, yet it still delivers pretty good performance. In fact, it outperformed the comparably priced Visiontek GeForce4 MX 440 in our tests.

The FDX 8500 LE uses ATI's Radeon 8500 LE video chip, a "lite" version of ATI's Radeon 8500 chip. The FDX 8500 LE has 64MB of DDR-SDRAM, which is typical of most budget video cards. The card only has a 250MHz RAMDAC, rather than the 350MHz RAMDAC most video cards have, but this doesn't



AP64D-H

\$240
Gigabyte
(626) 854-9338
www.giga-byte.com





3D Prophet FDX 8500 LE

\$145
Hercules
(877) 484-5536
(514) 279-9960
www.hercules.com



I knew the
FDX 8500 LE
wouldn't keep
pace with
most of the
video cards
in this roundup,
but it performed
pretty well
considering
how little
it costs.

reduce the FDX 8500 LE's performance. The card is compatible with most major OSes, including Win98/Me/2000/XP.

I knew the FDX 8500 LE wouldn't keep pace with most of the video cards in this roundup, but it performed pretty well considering how little it costs. The FDX 8500 LE's 3DMark2001 score of 7,744 was more than 2,000 points higher than the Visiontek GeForce4 MX 440's score, even though that card costs as much as the FDX 8500 LE.

Despite its considerably lower 3DMark2001 scores than the GeForce4 Ti 4600 and Radeon 8500 video cards, the FDX 8500 LE did manage to keep up (for the most part) with those pricier video cards in the Quake III test. The FDX 8500 LE's frame rate at 800 x 600 was 190.2fps. This was the slowest frame rate at this resolution among the cards we tested, but it's still quite good for Quake III. At 1,024 x 768, the FDX 8500 LE's frame rate was 172.4fps, surpassing the Visiontek GeForce4 MX 440's frame rate of 157.2fps. When we increased the resolution to 1,600 x 1,200, the FDX 8500 LE's frame rate dropped to 89.4fps, which is still adequate to run Quake III well.

The Serious Sam demos took a toll on the FDX 8500 LE. The video card's frame rate in the Karnak demo was only 51.2fps. The frame rate in the Metropolis demo was slightly better at 59.1fps. Both rates were considerably better than those of the Visiontek GeForce4 MX 440, but the frame rates are a little too low for comfort. We noticed some slowdowns in gameplay when running Serious Sam with this video card, although they weren't especially disruptive.

I'd choose the Hercules 3D Prophet FDX 8500 LE over a GeForce4 MX 440 video card. The FDX 8500 LE's price is almost identical to a GeForce4 MX 440, yet the FDX 8500 LE clearly performs better. The FDX 8500 LE has its limits, but it's a good card for those on a budget.

Leadtek WinFast A250 Ultra TD

If you measure the value of a 3-D card strictly by the frame rates it can produce, a GeForce4 Ti 4600-based video card is very valuable indeed (this is reflected in the sticker price, as well). It seems NVIDIA wows us every year with an amazingly powerful new GPU that blazes through video game graphics. The NVIDIA Ti 4600 is this year's king of the GPUs (at least for now), and Leadtek's WinFast A250 Ultra TD video card is the king's primary ambassador.

The WinFast A250 Ultra TD sports 128MB of DDR-SDRAM and has a wide range of refresh rates, from 60Hz to 240Hz.

Head To Head

	ATI Radeon 8500 128MB
Price	\$299
Interface	AGP
Chipset	Radeon 8500
RAM	128MB DDR-SDRAM
Max.Color Depth(bits)	32-bit
Max.Resolution	2,048 x 1,536
Available Refresh Rates	85Hz to 200Hz
RAMDAC (MHz)	400MHz
APIs Supported	OpenGL, DirectX
Operating Systems	Win98/Me/2000/XP
3DMark2001	8,496
Quake III	
Frames per second @	
800 x 600	198.6
Frames per second @	
1,024 x 768	186.0
Frames per second @	
1,600 x 1,200	108.4
Serious Sam	
Karnak Coop	66.25
Metropolis Coop	75.9
CPU Rating	2.5
Manufacturer	ATI
Final Word	Respectable performance for the price
URL	www.ati.com

Like the other video cards in this review, the WinFast A250 Ultra TD supports Win98/Me/2000/XP.

In our 3-D tests, the card ruled the roost, posting the best 3DMark2001 score at 9,652. That's far better than most of the video cards we tested for this roundup and is the highest 3DMark2001 score we've seen to date, period.

The WinFast A250 Ultra TD also dominated the competition in the Quake III demo and the Serious Sam demos. In our Quake III tests, the WinFast A250 Ultra TD's frame rate was 208.2fps at 800 x 600. At 1,024 x 768, the frame rate was 204.3fps. Both frame rates were the fastest among the video cards we reviewed. The Visiontek Xtasy GeForce4 Ti 4600 barely edged out the WinFast A250 Ultra TD for the fastest frame rate at 1,600 x 1,200, with the WinFast card running 157.7fps to the Visiontek's 158.2fps. Still, the WinFast A250 Ultra TD showed a lot of muscle in the Quake III test.

The A250 Ultra TD continued its dominance in the Serious Sam tests. The video card's frame rates in the Karnak (84.4fps) and Metropolis (100.8fps) demos were the highest among all the cards we reviewed. We were impressed, to say the least.

The WinFast A250 Ultra TD is tough to beat when it comes to raw power, but it's not without an Achilles' heel, namely its price. You'll have to pay more than \$400 for the privilege of having the fastest video card in your neighborhood.

Visiontek Xtasy GeForce4 MX 440

The Xtasy GeForce4 Ti 4600 receives my well-earned kudos, but I can't do the same for the Xtasy GeForce4 MX 440. The MX 440 is a



WinFast A250 Ultra TD

\$400

Leadtek

(510) 490-8076

www.leadtek.com



Gigabyte AP64D-H	Hercules 3D Prophet FDX 8500 LE	Leadtek WinFast A250 Ultra TD	Visiontek Xtasy GeForce4 MX 440	Visiontek Xtasy GeForce4 Ti 4600
\$240	\$145	\$400	\$140	\$380
AGP	AGP	AGP	AGP	AGP
Radeon 8500	Radeon 8500LE	GeForce4 4600	GeForce4 MX440	GeForce4 4600
64MB DDR-SDRAM	64MB DDR-SDRAM	128MB DDR-SDRAM	64MB DDR-SDRAM	128MB DDR-SDRAM
32-bit	32-bit	32-bit	32-bit	32-bit
2,048 x 1,536	2,048 x 1,536	2,048 x 1,536	2,048 x 1,536	2,048 x 1,536
Up to 200Hz	Up to 200Hz	60Hz to 240Hz	60Hz to 240Hz	60Hz to 240Hz
NA	250MHz	350MHz	dual 350MHz	350MHz
OpenGL, DirectX	OpenGL, DirectX	OpenGL, DirectX	OpenGL, DirectX	OpenGL, DirectX
Win98/Me/2000/XP	Win98/Me/2000/XP	Win98/Me/2000/XP	Win95OSR2/98/NT 4.0/Me/2000/XP	Win95/98/NT 4.0/Me/2000/XP
8,169	7,744	9,652	5,517	9,616
197.4	190.2	208.2	192.8	207.4
183.3	172.4	204.3	157.2	203.5
103.8	89.4	157.7	74.4	158.2
63.9	51.2	84.4	38.5	78.7
70.8	59.1	100.8	41.5	92.3
2.5	2.5	4	2	3.5
Gigabyte	Hercules	Leadtek	Visiontek	Visiontek
Decent alternative to NVIDIA-based cards	Excellent budget video card	Fastest graphics in town	Can't keep up with comparably priced video cards	Great performance
www.giga-byte.com	www.hercules.com	www.leadtek.com	www.visiontek.com	www.visiontek.com

The
Serious Sam
demos further
underscored
the Xtasy
GeForce4
MX 440's
lack of
competitive
spirit.

Xtasy GeForce4 MX 440

\$140

Visiontek

(800) 726-9695

(847) 360-7500

www.visiontek.com



video card for 3-D gamers on a budget, similar to the Hercules 3D Prophet FDX 8500 LE, but the MX 440's performance wasn't that impressive. I knew it wouldn't be on the same plateau as a video card using a GeForce4 Ti 4600 or ATI Radeon 8500 GPU, of course, but I still expected more.

The Xtasy GeForce4 MX 440 uses 64MB of DDR-SDRAM and has refresh rates from 60Hz to 240Hz. Like the Xtasy GeForce4 Ti 4600, it supports numerous OSes, including Win95/OSR2/98/NT 4.0/Me/2000/XP.

The Xtasy GeForce4 MX 440 didn't get off to a good start in our 3-D tests, bombing 3DMark2001 with a total score of only 5,517, the lowest (by more than 2,000 points) in this roundup. The results weren't much better in our other tests.

Although the MX 440 maintained respectable frame rates in the Quake III tests, the frame rates were considerably slower than the other video cards. The exception was at 800 x 600, where the Xtasy GeForce4 MX 440's frame rate was 192.8fps. That's slightly faster than the Hercules FDX 8500 LE's frame rate. However, the MX 440 dropped to last place at 1,024 x 768 (157.2fps) and at 1,600 x 1,200 (74.4fps). These frame rates are good enough to ensure smooth gameplay in Quake III, but they're not competitive with the other video cards we reviewed.

The Serious Sam demos further underscored the Xtasy GeForce4 MX 440's lack of competitive spirit. The video card's frame rate during the Karnak demo was only 38.5fps, which is

barely acceptable. The card did better in the Metropolis demo with a frame rate of 41.5fps, although that's still pretty weak.

Many PC gamers would like to upgrade the power of their PC's graphics without spending oodles of money, and the Xtasy GeForce4 MX 440 should be just the video card for them. However, those users would be better off with a Hercules 3D Prophet FDX 8500 LE, which costs about the same as the Xtasy GeForce4 MX 440 but delivers better performance. Overall, we weren't very impressed with the GeForce4 MX 440.

Visiontek Xtasy GeForce4 Ti 4600

Visiontek's Xtasy GeForce4 Ti 4600 was nipping at the heels of the Leadtek WinFast A250 Ultra TD in nearly every one of our tests. The difference in performance between the two cards is negligible for the most part, and the specifications for both cards are nearly identical.

Like the WinFast A250 Ultra TD, the Xtasy GeForce4 Ti 4600 uses NVIDIA's new GeForce4 Ti 4600 GPU and has 128MB of DDR-SDRAM. The Xtasy GeForce4 Ti 4600 also has refresh rates from 60Hz to 240Hz like the WinFast card. The Xtasy is compatible with a wider range of OSes, however, including Win95/98/NT 4.0/Me/2000/XP.

The Xtasy GeForce4 Ti 4600 was solid in every phase of testing. Its 3DMark2001 score of 9,616 was just slightly behind the WinFast A250 Ultra TD's top 3DMark2001 score and considerably better than the other video cards in this roundup.

The Xtasy GeForce4 Ti 4600 kept up with the WinFast video card through almost all of the 3-D tests. At an 800 x 600 resolution, the Xtasy's frame rate was 207.4fps, only 1fps less than the WinFast card. At 1,024 x 768, the Xtasy's frame rate was 203.5fps, again only about 1fps slower than the WinFast's score. The Xtasy GeForce4 Ti 4600 surpassed the WinFast at 1,600 x 1,200, posting a frame rate of 158.2fps to the WinFast's 157.7fps.

The Serious Sam tests didn't go quite as well for the Xtasy GeForce4 Ti 4600 as for the WinFast video card. The Xtasy's frame rate of 78.7fps in the Karnak demo is very good to be sure, but it was several frames per second slower than the WinFast's frame rate.

The results were similar in the Metropolis demo. The Xtasy GeForce4 Ti 4600's frame rate was 92.3fps in Metropolis, which was faster than every other video card I reviewed, with the exception of the WinFast A250 video card.

The face-off between the Xtasy GeForce4 Ti 4600 and the Leadtek WinFast A250 Ultra TD ended in almost a draw. The WinFast card is slightly more powerful than the Xtasy GeForce4 Ti 4600 is, but the Xtasy costs about \$xx less.

The Card For You

I've shown you a variety of video cards at different price classes so you can determine how to get the most bang for the buck.

Leadtek's A250 Ultra TD is my pick of the cards in this review, but I also liked Visiontek's Xtasy GeForce4 Ti 4600. Although the ATI-powered video cards performed well, they couldn't quite hang with the GeForce4 Ti 4600 video cards in the end. I'd rank the Hercules 3D Prophet FDX 8500 LE as the best choice for power users who are on a budget. ▲



Xtasy GeForce4 Ti 4600

\$380

Visiontek

(800) 726-9695

(847) 360-7500

www.visiontek.com



by Michael Sweet

Two More For Your Consideration

3Dlabs Wildcat II 5000

All consumer 2-D/3-D video cards support OpenGL, the API that makes fast, luscious 3-D games possible. But some 3-D cards use OpenGL for other purposes, such as for designing the 3-D shooters we love so much. Video cards like the Wildcat II 5000 from 3Dlabs use OpenGL to run intense design



3Dlabs Wildcat II 5000

\$1,199

3Dlabs

(800) 464-3348

(408) 530-4700

www.3dlabs.com



apps, such as CAD programs, to create anything from new buildings to lifelike animated characters. These cards are designed much differently than consumer 2-D/3-D cards.

The Wildcat II 5000 has multiple video chips and 80MB of RAM; 32MB of RAM is reserved solely for the frame buffer, with another 32MB reserved for the texture buffer. One of the card's more striking features is its size, which measures about twice as long as an average video card.

The 5000 has a RAMDAC of 300MHz and a maximum resolution of 2,048 x 1,152. Its refresh rate tops out at only 90MHz, but that's still sufficient to prevent screen flicker. The Wildcat II 5000 is only compatible with WinNT and Win2000, which is a good hint that this isn't your everyday video card for home users. Here's another clue: The 5000 costs as much as \$1,200.

Matrox Millennium G550

Most video cards you see at CompUSA or Best Buy are 2-D/3-D cards equally adept at running games and office programs. These double-duty video cards might be just the ticket for

many users, but plenty of users don't need or want a 3-D video card at all. If you work in a small office or home office, chances are you aren't (or at least shouldn't be) playing video games all day, so why spring for a \$300 2-D/3-D video card? An office PC needs to be good at accelerating 2-D graphics, however, and many office users are discovering the benefits of running multiple displays from one system. Most 2-D/3-D video cards handle 2-D graphics pretty well, but the Matrox Millennium G550 video card is a specialist in 2-D graphics. It

accelerates 2-D applications better than other video cards I've tested and also has excellent multiple display capabilities.

The Millennium G550 uses a Matrox G550 video chip and has 32MB of DDR-SDRAM. The card's RAMDAC of 360Hz is slightly faster than most. It also supports more OSes than other video cards. You can pop a Millennium G550 into a PC running Win98/NT4.0/Me/2000/XP or Linux. The Millennium G550's refresh rates are from 85Hz to 200Hz.

The Matrox Millennium G550 produced the highest SYSmark2002 scores of any video card we've tested so far, which is a testament to this card's 2-D acceleration. The card's Overall 155 score was based on an impressive Office Productivity score of 100 and a solid Content Creation score of 239. In addition, the card only costs \$125. The Matrox Millennium G550 isn't for everyone, but if you need a good 2-D card without the 3-D eye candy, the Millennium G550 is the way to go. ▲



Matrox Millennium G550

\$125

Matrox

(800) 361-1408

(514) 822-6000

www.matrox.com





**RARELY DO YOU GET TO USE THE WORDS "INGENIOUS"
AND "MUSCLE-BOUND" IN THE SAME SENTENCE.**

Introducing the all-new Chevy Avalanche with its exclusive Midgate.[™] The only vehicle that changes from an SUV to a

*Excludes other GM vehicles. Chevy, Avalanche and Midgate are trademarks of the GM Corp. ©2001 GM Corp. Buckle up, America!

PROUD SPONSOR OF
THE 2002 U.S. OLYMPIC TEAM



full-size pickup. chevy.com/avalanche

AVALANCHE



LIKE A ROCK



ClearPC

\$199
ClearPC
(250) 721-3444
www.clearpc.ca



ClearPC

Back in the hair-spray days of the 1980s, Swatch came out with transparent watches. Other companies have since caught on to the phenomenon, releasing transparent cell phones, engine covers, knickers, and Macs. There's just something satisfyingly useless about buying products and being able to look inside while you use them, isn't there? Sure it's pointless but a) it looks cool and your pals will be dead impressed and b) you get to see some moving parts. Enter the ClearPC case, which case modders can customize to their hearts' content. At 13 pounds, the case is also useful for schlepping to LAN parties. And yes, of course you get to look at LEDs blinking, fans whizzing, and so on. Cool.

The ATX chassis, which is pretty standard-sized for a midtower at 16-13/16 inches high x 8-11/16 inches wide x 18-3/8 inches deep, is made from 90% 3/16-inch cast clear Lucite. A screwless side panel allows for easy access to hardware for upgrades. The front bezel is also removable, and the entire unit is held together with

finger-tightened Cap nuts. There is a pre-cut 80mm intake and 80mm exhaust for cooling purposes, which probably isn't enough for hardcore overclockers but ample for most users nonetheless. ClearPC also sells "clear" 80mm case fans for an even fuller, or "clearer," effect. Each case is handmade, thus all corners and edges have been rounded off.

Installation of hardware is a little more tricky than normal and requires a bit more time, patience, and even sticky tape. (Mainboard screws need to be taped to the back of the chassis.) Media drives are mounted via plexi tubes to give a nifty "suspended" look. There are enough parts to set up one HD, one media drive, and a floppy, so you'll need to buy extras for beefier PCs.

Think of this case if you are a fashion fiend and willing to part with \$199.99. Expect to spend even more on accessories, as well. But since when has looking good been cheap? ▲

by Alex "Sharky" Ross



ATC-210 VX2

\$250
Cooler Master
(510) 770-8566
www.coolermaster.com



Cooler Master ATC-210 VX2

Looks, weight, and cooling have become much more of a given since customizing and overclocking hit the mainstream. We geeks might dress like crap in the real world, but we sure know how to gussy up our becursed PCs in sexy ATX cases. The advent of extremely high-frequency CPUs, graphics chips, and memory has helped to rid the industry of those big/boring/beige cases. Thanks to Cooler Master, aluminum cases have really caught on, with the ATC-201 becoming an overclocker's expensive favorite.

The newer ATC-210 is structurally similar to the ATC-201 (see page 34 of the January 2002 *CPU*) but targeted at a slightly different audience. Where the ATC-201 is meant for extreme cooling, the ATC-210, with only dual 80mm exhaust fans (one at the top and one at the rear), is aimed at users more interested in looks. The ATC-210's acrylic front panel doesn't allow for Cooler Master's intake-to-exhaust method as on the ATC-201, and temps are liable to be slightly higher. Still, the positioning of the fans is excellent; the design allows for good airflow paths.

Made from aluminum alloy, the case has four 5.25-inch bays, two 3.5-inch bays, and room enough for four hard drives. The front panel has dual USB 1.1 ports. Most of the internal pieces, such as the motherboard tray, HD/CD-ROM bays, and even the power supply cage, are also made from aluminum. The case is accessible, thanks to "tool-less" assembly/disassembly thumb screws, and the pull-out motherboard tray makes installation and upgrades speedy. The tray sports two stability bars and glides out squeaky clean.

Perhaps the biggest selling point is the front acrylic translucent bezel, which comes in either fog blue or fog green. Thanks to two magnets, the Plexiglas door opens and shuts, giving you easy access to your drives and, in truth, does give the case a rather distinctive look. The ATC-210 is one of the classiest hybrid PC cases on the market. It not only looks fab, but at under 13 pounds, it's also lightweight. Need I mention, though, that it's quite pricey at \$250? ▲

by Alex "Sharky" Ross

Intel Pentium 4 2.4GHz

With Intel's superb 0.13-micron manufacturing process, higher clock speeds are not much of a problem, hence the release of the P4 2.4GHz.

Built with newer 300mm wafers and using the same Northwood core as you'll find on the 2.2GHz and 2A, the P4 2.4GHz features 512KB of on-die L2 cache. Intel will shortly be releasing P4s running at 133MHz (quad-pumped 533MHz), which should give it some more juice. For the time being, the P4 2.4GHz runs with a 100MHz FSB and uses a multiplier of 14x.

The deep pipeline and 0.13-micron process of the P4 proves its worth yet again where over-clocking is concerned. If 2.4GHz seems slow, you'll be happy to know overclocking to 3GHz is a simple affair that doesn't require any fancy cooling techniques. (Remember the days of the Celeron 300A?) If overclocking is your thing, and you're willing to spend so much money on a CPU, you'll undoubtedly have some fun.

AMD Athlon XP 2100+

Intel's Pentium 4 2.2GHz vs. AMD's Athlon XP 2000+ battle in the March CPU was a close one, close enough to make AMD up the ante once more by releasing yet another Athlon XP, the 2100+. AMD's model numbers portray overall performance figures rather than actual clock speed, and when compared to Intel's P4 frequency, truth be told, the numbers are actually quite modest. The 2100+ translates into 1.73GHz with a multiplier of 13x. The 2100+ is 67MHz faster than the 2000+. 4% isn't exactly a quantum leap, is it? Hence the performance delta between the 2000+ and the 2100+ is obviously negligible but just enough to tip the scales against the P4 2.2GHz.

The 2100+ retains the 0.18-micron Palomino core from previous Athlon XPs and includes 128KB of L1 cache and 256KB of full-speed, on-die L2 cache. In line with other Athlon XPs, the 2100+ uses a 1.75V core voltage and a 266MHz frontside bus.

If you already own an Athlon XP, you're in great shape, and there's absolutely no need

It's been quite awhile since Intel has had a performance lead over AMD, but with a huge clock speed advantage, the P4 2.4GHz takes the speed king crown this time around. With a price of \$562, though, it's going to take some more work on Intel's part to get the pricing down before the majority of gamers/enthusiasts decide to have "Intel inside" again. The Athlon XP range and DDR platform still offer the best bang for your buck if you are a gamer and concerned about better IPC (instructions per clock cycle) and loathe Rambus. But if you are looking at upgrading to a P4 system, the enthusiast in me is going to recommend you wait just a wincey bit longer for 533MHz FSB to turn up. ▲

by Alex "Sharky" Ross

Intel & AMD Face Off

The P4 2.4GHz wins most of the benchmarks tally by a small margin. Quake III engine-based games with the P4 and RDRAM do stand out, but you should take into account the average pricing when buying the entire platform. Obviously the Athlon XP is kinder to your bank account; you be the judge of just how much a few extra frames per second are worth to you.


to upgrade to a slightly faster stepping. After all, the \$420 isn't exactly cheap. Instead, sit tight for the 0.13-micron Thoroughbred core. If you're living life in the truck lane with a slower processor and just can't wait for the Thoroughbred core, then look at better value-priced Athlon XPs. Currently there are some great deals on the 1800+, giving you a much better bang for the buck.

Perhaps the 2100+ ends up being more of a political statement as AMD has rounded off things with the final Palomino core-based CPU. Of course, politics can get zesty at the best of times, hence Intel has responded swiftly by hitting back with the Pentium 4 2.4GHz, which just beats the 2100+. To quote C-3P0 (that human-cyborg relations chap from "Star Wars"), "Will it never end?" ▲

by Alex "Sharky" Ross


Pentium 4 2.4GHz

\$562
Intel
(408) 765-8080
(916) 377-7000
www.intel.com



	P4 2.2GHz	P4 2.4GHz	XP 2000+	XP 2100+
PCMark 2002 CPU	5387	5888	5104	5003
SYSmark2001	221	225	189	190
3DMark 2001	8326	8527	8446	8537
SYSmark2001	221	225	189	190
Comanche 4	41fps	45fps	41fps	42fps
Jedi Knight II	95fps	101fps	93fps	90fps
Quake III	272fps	289fps	250fps	255fps
Serious Sam	115.2fps	128.2fps	133.2fps	141.2fps
Price (CPU, memory and mobo)	Approx. \$805	\$910	\$420	\$480

*Tested on a WinXP Pro system with a GeForce 3 Ti 500



Athlon XP 2100+

\$420
AMD
(800) 538-8450
(408) 732-2400
www.amd.com



NVIDIA GeForce4 Ti 4200 vs. ATI Radeon 8500LE On The Cheap

GeForce4 Ti 4200

\$179
NVIDIA
(408) 486-2000
www.nvidia.com



GeForce4 Ti 4200 vs. Radeon 8500LE

	GeForce4 Ti 4200	Radeon 8500LE
3DMark 2001SE	9327	8052
3DMark 2001 (Nature)	38.8	40
Quake III	228fps	196.5fps
Quake III 2XAA	163fps	91.3fps
Comanche 4	43.2fps	36.6fps
Serious Sam TSE	142fps	118fps

**Tested with a P4 2.2GHz with 256MB PC800 at 1,024 x 768 x 32

Radeon 8500LE

\$199
ATI
(905) 882-2600
www.ati.com



NVIDIA's GeForce4 Ti 4600 is the current undisputed 3-D champ. ATI's Radeon 8500 isn't too bad, either. Both rip through complex 3-D engines without breaking into too much of a sweat. But there's a catch: How many of you actually fancy dropping \$400 on a board? It doesn't help knowing that in six months, both NVIDIA and ATI will most likely release a speedier chip. So what are frustrated 3-D gamers with no more than \$200 burning a hole in their pocket to do? Return that XM satellite car radio? No, Obi Wan, there is another hope . . .

Ti 4200. NVIDIA's segmentation strategy allows the company to sell another GeForce4 Ti product (with the NV25 chip) at a sub-\$200 price point, which means gamers can ignore the GF4 MX in favor of a new and more exciting part called the Ti 4200. Priced at \$179, the board resembles a GF4 MX in terms of layout, with 64MB of 500MHz DDR-SDRAM rated at 4ns (no BGA stuff here, sorry) and a core frequency of 250MHz, but don't let that fool you.

Like other GF4 Tis, the Ti 4200 harbors an NV25 chip with Accu-view Antialiasing, an nfiniteFX II engine, Lightspeed Memory Architecture II, and nView capabilities. Unlike the GF4 MX, the Ti 4200 sports Pixel Shaders and dual Vertex Shaders and hence is a proper DX8-class product. A DVI-out came with my reference board but will remain at the OEM's discretion; price will reflect that. Although the reference board featured a less fancy heat sink and fan combo than the Ti 4400 and 4600, the 4200 was nonetheless overclocking-friendly.

8500LE. ATI is slightly ahead of the game, having already had a \$199 (cheaper online) Radeon 8500LE product in stores for a couple of months along with an even cheaper 64MB version, which you can find for less than \$150. Indeed, the LE is largely responsible for NVIDIA having to counter with the Ti 4200. It seems as though the times are a changin' with NVIDIA having to be reactive as opposed to proactive.

Just like the Radeon 8500, the LE uses the same R200 core, featuring ATI's TrueForm, SmartShader, Hyper Z II, Pixel Tapestry II, Charisma Engine II, and SmoothVision technologies. The major hardware differences amount to a drop in memory frequency to 500MHz and core clock speed to 250MHz (25MHz slower than the 8500) and the removal of the DVI-out port for the LE. Such snipping helps make the LE cheap. The LE boasts 128MB of DDR memory, which is remarkable at such a low price point, especially when considering the current upsurges in memory prices. In fact, these upsurges may have forced NVIDIA to drop its original Ti 4200 spec from 128MB to 64MB. The peak bandwidth of the LE is an impressive 8GBps. But ATI's extra 64MB hardware advantage doesn't quite pull up the card's socks enough to beat out the Ti 4200, except in the 3DMark 2001SE Nature test, where complex scenes require heaps of texture memory. In every other test, NVIDIA wins the performance trophy.

The benchmarks have it. Both ATI and NVIDIA have done the right thing by releasing DX8-class products below \$200 in hopes of getting their technology to the masses. Now it's up to game developers to put their backs into pixel & vertex shader effects: Bring it on. Although either card is a real bargain, if I had to recommend one over the other (guess I do, actually), I'd let the benchmarks speak for themselves. Thus it comes down to the GF4 Ti 4200. If turning on Anti-aliasing is important to you, NVIDIA's Accu-view engine in the Ti 4200 seems to widen the performance delta even more significantly.

ATI's Radeon 8500LE is already on sale, with Hercules and Gigabyte to follow suit shortly. Expect current GF4 Ti OEMs to inject the Ti 4200 into their line-ups by the time you read this. Although the Ti 4600 is flashier, it's the Ti 4200 that is a far better bang for your buck and will likely end up in many more PCs. The rumor is that Ti 4200 prices will possibly drop down to \$150, which leaves you another \$29 to spend on wooly socks for Uncle Albert, who bought you that Athlon XP for Christmas. ▲

by Alex "Sharky" Ross

HP hn200w Wireless Gateway Router

Here's a step forward in 802.11b wireless networking. HP's hn200w has the easiest setup I've seen for a wireless gateway router. Its friendly installation wizard is just what novices have been waiting for: a way to set up a wireless network without asking anyone for help.

The attractive hn200w acts as an 11Mbps 802.11b access point, a four-port 10/100Mbps Ethernet switch with an uplink port, and a NAT firewall that even checks outbound traffic and TCP/IP ports. The router supports Win98/98SE/Me/XP/2000. The hn200w's quick start guide is minimal, but there's much more documentation on its installation CD.

Another boon for new networkers is that you can configure this router with an easy GUI utility, as well as a traditional browser-based one. The GUI is simply arranged with very understandable tabbed pages, while the browser-based menus offer more advanced settings.

The hn200w uses up to 128-bit WEP to secure 802.11b access. You'll have to approve

new users entering the WLAN area before they can access anything, which adds more security against interlopers. If you like, you can restrict a node's Internet access or kick it off the WLAN altogether.

One of the hn200w's coolest simplification tools is its application hosting. The GUI-style configuration utility lets you choose specific application and game titles from a lengthy, updatable list. The hn200w will then automatically configure itself for necessary access. You also can set permissions the usual way in the browser-based settings.

We connected the hn200w to a WinMe IBM PC300, several WinXP WinBook and IBM notebooks, and a DS-3 Internet connection. It sustained a reasonable network connection even 120 paces across an office building. This gateway is functional, easy to use, and has very nice features. ▲

by Marty Sems



hn200w Wireless Gateway Router

\$219.99

Hewlett-Packard

(800) 752-0900

(888) 999-4747

www.hp.com



SyNet USB Pack

OK, class: What happens when you try to "network" two PCs with a regular USB cable? Good answer, Timmy. Nothing good and possible damage.

USB wants to be FireWire in so many ways . . . and it just got closer with the SyNet USB Pack networking kit. It connects two PCs running Win98/98SE/Me/2000 Professional/XP but can support as many as 17 systems with proper equipment. Literally a 10-foot modified cable and a CD, the USB Pack might be all you need to transfer huge files, share an Internet connection, or share a printer.

Obviously, this kit is ideal for home users who want to network two computers but have an aversion to opening either's case to install NICs. It also makes a handy way to form an ad-hoc network between notebooks in a meeting without having to deal with security permissions on the company's existing network.

To implement the USB Pack, use the included CD to install drivers on both PCs.

The computer with the Internet connection will be the server, and the other will act as a client. Set up IP addresses and other settings according to the instructions, then install the Internet Sharing Program.

Of course, you shouldn't expect sky-high network speeds out of a USB 1.1 device. Still, SyNet says to expect transfers greater than 6Mbps. That's faster than a null modem cable over Direct Cable Connection, which I remember as being fast enough for some head-to-head Half-Life.

We connected an IBM PC300 running WinMe and a WinBook X2 running WinXP Home with the USB Pack. The SyNet cable sustained a 1.44Mbps (180KBps) download of a game demo from our DS-3 broadband connection. For a casual networking method, the USB Pack looks pretty impressive. ▲

by Marty Sems



USB Pack

\$49.99

SyNet

(866) 877-9638

(201) 945-3360

www.synet-router.com



Samsung SpinPoint SP8004H 80GB



SpinPoint SP8004H 80GB

\$129 (street price)

Samsung

(800) 726-7864

(201) 229-4000

www.samsunghdd.com



Although Samsung isn't about to unseat the top names in the hard drive industry in consumers' minds just yet, it is making bolder strides toward doing so. It may take a while, but drives such as the SP8004H could one day make you boast to a geek friend, "Yeah, it's a Samsung."

Samsung's new SpinPoint P40 series of hard drives combine 40GB platters with 7,200rpm spindle speeds. I'm pretty impressed with the two examples the company sent me. As I mentioned in my review of the 40GB SpinPoint SP4002H, both SpinPoints approach the Seagate Barracuda ATA IV in performance and shock tolerance ratings.

This 80GB version actually beat its 40GB little brother with higher average read/write rates and a lower average access time. We tested both on a 1.8GHz Pentium 4 system with WinXP Pro and 512MB Rambus memory. We also had to check the SP8004H's jumper settings a little more closely than usual, as WinXP

couldn't see more than the first 40GB until we rearranged them.

In WinBench, the SP8004H predictably mirrored the 40GB Samsung in our Business Disk (7,170KBps) and High-End Disk (24,600KBps) tests. However, the SP8004H pulled ahead of its brother in HD Tach, for the most part. It charted a 41.5MBps average read rate (43.7MBps maximum) and a 19.2MBps average read rate (21.9MBps maximum). Its 13.2ms random access time matched or beat Western Digital's 100GB and 120GB drives, which currently top our EIDE Top 5 list.

Like its sibling, this SpinPoint has a 2MB cache buffer, Ultra ATA/100 support, and a 31.5dB idle acoustic rating (33dB during reads and writes). The drive should last for 50,000 start/stop cycles, Samsung says, and it has a 350G nonoperating 2ms shock tolerance. And yes, that's 80GB for \$129. Incredible. ▲

by Marty Sems

IBM Deskstar 120GXP 120GB



Deskstar 120GXP 120GB

\$329

IBM

(888) 426-5214

(507) 253-4110

www.ibm.com/harddrive



Serial innovator IBM recently entered the 120GB weight class with the Deskstar 120GXP. It's a 7,200rpm challenge to Western Digital's 120GB WD1200BB and WD1200JB.

This EIDE hard drive enjoys several IBM drive hallmarks, including rigid glass disk platters and ramps that let the read/write heads rest off the platters when the drive is turned off. It also employs IBM's AFC (antiferromagnetically-coupled) recording media on each platter. However, its 29.7Gb per square inch areal density doesn't really show off how densely IBM predicts AFC will one day store data. As do other drive manufacturers besides Western Digital, IBM offers this drive in smaller capacities down to 20GB.

We benchmarked the 120GXP on a 1GHz PIII system with 384MB of SDRAM, WinMe, and a HighPoint DMA/100 adapter. Both its WinBench 99 scores struck near the top rank, with a 6,930KBps Business Disk rating and a 24,900KBps High-End Disk score. The 12GXP followed up its exceptional 12ms random-access time in HD Tach with terrific 38.4MBps

average read (48.4MBps max) and 24.7MBps average write (34.0MBps max) rates.

A low 8.5ms advertised read seek time and good management of its 2MB cache help this Deskstar fly. It appears that the 120GXP's lower access time and higher write rates put it above the WD1200BB with a 2MB cache (\$245), although it can't quite top the WD1200JB with 8MB cache (\$339). The 120GXP is a tougher and quieter drive than both, though. Its 350G nonoperating 2ms shock tolerance (400G in the 80GB or smaller versions) is outstanding, while its 31dB idle and 32dB to 37dB seek acoustical ratings are things to brag about.

IBM gives the Deskstar 120GXP a three-year warranty; I give it a big thumbs up. (NOTE: After recent concern over IBM's recommendation of 333 powered hours per month for the Deskstar 120GXP, the company stated that the drive can be and is used in 24/7 servers and that exceeding the 333-hours figure has no effect on the warranty.) ▲

by Marty Sems


COMPLETE DISASTER!



TOTAL RECOVERY

Just a moment ago things were going great. Your figures finally added up, that critical report was done and your MP3s were downloaded. That was then. Now you're looking at a blue screen on a dead machine.

But if you had Drive Image® 2002 you could sit back and relax knowing your entire system - including exact copies of irreplaceable Word™ documents, financials, E-mails, photographs, multimedia, your operating system, programs and preferences - was safely stored, and recovery is just a matter of reload and go!



With an easy-to-use Windows® interface, Drive Image 2002 allows you to quickly make a complete copy of your entire system onto a Zip®, Jaz®, CD-R/W, network drive, disk partition or any other suitable storage device. So, when disaster strikes, all your valuable data is ready and waiting for you to put it right back where you left it.

For hassle-free data and system recovery right now, either visit your local software reseller, go online to www.powerquest.com/smart1 or call 1-888-497-9998 and order Drive Image today.



Fry's ELECTRONICS

Office DEPOT



STAPLES

COMPUSA



Digiset DUO-64



DUO-64

\$229

Digiset

(973) 455-7084

(866) 434-4473

www.digiset.com



Most people love the idea of getting a two-for-one deal, and that's just what you get with the Digiset DUO-64 digital audio player. It's a digital audio player that looks and acts like an audiocassette.

The DUO-64 has 64MB of on-board memory, but if that's not enough, you can buy a SanDisk MMC, which come in 32MB and 64MB versions. Or, you can get the DUO-96 with 96MB of on-board memory. The DUO-64, which is USB compatible, supports MP3 and WMA formats. Using the WMA format, the DUO-64 played for about 2:12 (hours:minutes). The player comes with two NiMH rechargeable batteries.

In addition to playing music and books from Audible.com, the DUO-64 can record voice clips. Just switch it to record mode and push the Play button. You can store voice recordings alongside other audio files.

As with all digital audio players, you have to download before you listen. But when I

tried to download music using Real Jukebox as the owners manual suggested, I found that the plug-in for Real Jukebox was missing. According to Digiset, this plug-in has been put on hold, so the company recommended using Windows Media Player. Digiset says it may add an announcement to its Web page informing customers of the change, but as of press time, Digiset had yet to post such an announcement on its site.

I found the missing plug-in annoying; something else that irked me was that the player didn't have a built-in LCD. Plus, the On/Off switch and battery cover were awkward to handle. Despite these factors, the DUO-64 is worth checking into if you're fond of unique gadgets. The DUO-64 plays quality digital music and gives you a little something extra with which to have fun. ▲

by Jennifer Seeman

Lite-On LTR-40125S CD-RW 40X/12X/48X



LTR-40125S CD-RW 40X/12X/48X

\$170 (street price)

Lite-On

(408) 935-5353

www.liteonit.com



CD-RW drives are reaching theoretical burning rates that almost sound like the read rates of CD-ROM drives. Leaving aside the little-seen Zen drives with 72X drive ratings, most commodity CD-ROMs top out at about 52X. It's hard to believe, but Lite-On and others are now selling 40X CD-R burners. Although this particular drive doesn't make the quantum leap in performance you might expect, it nevertheless took top honors as the fastest CD-RW thus far in our tests.

The most important test time for a CD-RW is how fast it burns a CD-R. The LTR-40125S burned a 427MB folder to an appropriately rated blank disc in 3:07 (minutes:seconds), a new record around here. The 5:15 it took to write the test folder to CD-RW is also pretty good.

We ran further tests with CD Tach 98 2.0 on a system with a 1GHz PIII, 384MB of PC100 SDRAM, a HighPoint DMA/100 controller, and WinMe. The highlights were an average 5,248KBps read rate, a 30.8X

weighted average drive rating, and a 6,956KBps maximum read rate (about 46.4X). Random and full-stroke access times were on the sunny side of ordinary at 78ms and 167ms, respectively.

The other main attraction to this zippy drive, besides SMART-BURN buffer under-run protection, is its support for the Mt. Rainier rewritable format (sometimes called CD-MRW). This interesting new CD-RW format, slated for support by WinXP and Linux, is supposed to enjoy drag-and-drop file writing ability from the OS. There's much more to Mt. Rainier, such as automatic formatting of new media and defective sector management, so it's worth looking into.

In short, this drive realizes the 40X rating's promise of increased speed better than Benq's 4012P, but there still seems to be room for improvement. It's a good investment. ▲

by Marty Sems

Adaptec XHub 4 Plus USB 2.0 Hub

I can honestly say this is the first review I've ever done of a picture frame. No, it's not an LCD; that was the first question on everyone's lips around here. A changeable—or continually changing—photo display would be way cool but would probably cost as much as a PDA or flat-panel monitor. This is just a regular photo frame. Oh yeah, it's also a four-port USB 2.0 hub.

Adaptec must have figured that as long as you're putting a hub on your desk within easy reach to plug and unplug things, it might as well double as a photo frame. The XHub 4 Plus' 1.875-inch square aluminum matte is a little difficult to pop off, but it's attractive and does the job. A thin Plexiglas window protects your photo from scratches.

The XHub 4 Plus' sturdy wire stand is actually functional. It corrals your USB cables

into a bunch, keeping their unruly selves near the unit's center of gravity to preclude tipping. The hub's four ports lie along its bottom edge. As is any good USB 2.0 hub, it's backward compatible with your USB 1.1 devices. It worked just fine for us between an HP ZT1170 notebook and a Canon S900 printer. A seven-port model costs \$129.99.

XHubs are stackable, although you'll have to remove any underlying Plus unit's photo frames first. You can daisy chain a maximum of five of these together using their uplink ports and included cables.

This somewhat pricey hub is nevertheless very practical. It comes with a USB 2.0 A-B cable, a power adapter, and a 1-year warranty. Photo not included. ▲

by Marty Sems



**XHub 4 Plus
USB 2.0 Hub**

\$79.99

Adaptec

(408) 945-8600

www.adaptec.com



Sony CLIE PEG-T415

If beauty is in the eye of the beholder, Sony's CLIE PEG-T415 has a nice body but not a very pretty face. The PEG-T415 has an eye-catching design, with a silver aluminum case and a flip cover that gives it a Star Trek-communicator feel. Scottie won't be able to beam you up, but you can beam data to other PDAs via a built-in infrared port (assuming the other PDA also has infrared abilities). The PEG-T415 is the essence of portability, weighing only 4 ounces and measuring 4.8 inches high x 2.9 inches wide x 0.4 inches deep.


The PEG-T415 has a handy jog dial and a Back button built into the upper-left side of the case. On the front are shortcut buttons for the calendar, address book, to-do list, and memo programs. The scroll button is in the middle and is functional, but I wish it were a little bigger.

The PEG-T415's biggest drawback is its screen. It has a relatively high 320 x 320 resolution but seems either too bright or too dark, depending on the angle. The PDA's backlighting feature slightly alleviates this problem, but not enough.

The PEG-T415 has 8MB of RAM, a 33MHz processor, and a Memory Stick slot to quickly add up to 128MB of RAM. The PDA uses the Palm 4.1 OS and comes with several programs, including Documents To Go, with which you can open Microsoft Word and Excel documents. I like the CLIE Remote Commander program, which lets you use the PEG-T415 as a remote for your television, VCR, or DVD player.

With a better screen, I'd recommend the CLIE PEG-T415 to anyone, but the PDA isn't a very useful product if you have to strain your eyes to see what you're doing. ▲

by Michael Sweet



CLIE PEG-T415

\$219.99

Sony Electronics

(888) 595-8246

(941) 768-7669

www.sonymstyle.com/clie



Logitech TypeAway



TypeAway

\$79.95

Logitech

(800) 231-7717

(510) 795-8500

www.logitech.com



Logitech's TypeAway keyboard is the new accessory all the cool Palm users will most likely be playing with soon. I had the pleasure of testing an engineering sample, but by the time you read this, you should be able to purchase the TypeAway (online, at least) for \$79.95.

When closed, the TypeAway is about the size of a cigarette case; it measures about 5.5 inches long x 3.5 inches wide x 0.4 inches deep and weighs 5.25 ounces. The case exterior is steel gray and durable, with a blue incandescent Logitech logo in the center. Press an inset button on the side of the end of the enclosure, and it opens into a slim keyboard with black keys, about 11 inches long.

The TypeAway is compatible with the Palm m125, m500, m505, and m515 series of PDAs. To use the TypeAway, you attach a Palm to the keyboard's hinged cradle; pull out a small support bar shaped like a squared-off "U," and you can prop up the Palm at a 45-degree angle.

For the Palm to recognize the keyboard, you must install the TypeAway software from a CD that bundles with the keyboard. This requires you to install the software on your PC (which must be a Mac or an IBM-compatible PC running Windows 95/98/NT/2000/Me/XP) and perform a HotSync to install the software on the Palm. The CD also includes WordSmith memo pad software for Palm and PC.

I liked the TypeAway a little better than Logitech's KeyCase, which is \$20 more but does provide unbeatable flexibility. (What other keyboard can you wrap around your Palm?) The 63-key TypeAway provides good tactile feedback and includes dual-function keys you can use to launch typical Palm functions or edit text. If you can imagine Q supplying James Bond with a Palm keyboard, the TypeAway would be it. ▲

by Cal Clinchard

Digital Dream l'espion



l'espion

\$57

Digital Dream

00 44 1304 248200 (UK)

www.digitaldreamco.com



If you ever caught the cartoon, you know that the tiny Smurfs lived in wondrous, happy Smurfland, tucked away in the forest and far from human eyes. Had the Smurfs conceived digital camera technology, they would've fashioned something like l'espion. l'espion (translated: "The Spy") comes from Digital Dream, based in Kent, England, which, judging by the weensy size of this camera, must be close to Smurfland.

This camera is the smallest I've seen. It weighs only 1.2 ounces with its single AAA battery installed, and its dimensions are about one-third the size of a standard cassette tape. l'espion comes with two buttons: One cycles through options such as picture quality and the self-timer, while the other functions as a shutter button. The only other moving part is a slider that controls the retractable viewfinder.

The camera, which captures photos using a 0.1MP CMOS sensor, stores photos using

2MB internal memory; you transfer them to your PC via the unobtrusive USB port. You can keep 20 photos at top resolution (352 x 288 dpi) or 80 at low resolution (176 x 144 dpi). Take photos singly or in a continuous string, or use l'espion as a 16fps Web cam.

The fixed lens focuses from about 39 inches to infinity.

Obviously, those hardware specs don't exactly boggle the mind (unless it's a very teeny Smurf mind), but this camera takes surprisingly good pictures. Under fluorescent lighting, pictures were clear and lush; outdoor photos in bright light were barely overexposed. Close-up shots were good, as well, even when I crept in closer than 39 inches. I did note some obvious barrel distortion, but it's not enough to ruin images.

"The Spy" is a cheap, immensely fun, and easy-to-use camera for both kids and adults. ▲

by Nathan Chandler

Apple iMac with SuperDrive

As I write this review, Apple is shipping more than 5,000 new iMacs every day, so the latest iMac revolution is underway. Of course, we won't know for sure whether Fate is smiling on the current iMac until it starts showing up in prime time on the WB, but things are looking good for Apple and its followers.

Now that a few mainstream publications have covered the iMac specs, I'm in the enviable position of being able to test the hype. Then again, there's not much else I can test because we don't have a Mac-compatible benchmarking utility. So make of it what you will; you'll find I enjoyed the iMac immensely but managed to find a few things worth griping about.

Specifications. For this review I looked at the iMac with a built-in Apple SuperDrive, which is a DVD-R/CD-RW combo drive. The drive reads DVD-ROMs at 6X, writes to DVD-R at 2X; for CDs the drive has 8X/4X/24X write/rewrite/read speeds. Looking for a 3.5-inch floppy drive? Look elsewhere; like previous models, this iMac has none.

The processor is an 800MHz PowerPC G4, which has a 100MHz bus speed and a 256MB L2 cache. The system has 256MB of SDRAM and a 60GB Ultra ATA hard drive. The processor, RAM, hard drive, and optical drive are the chief price determiners for the three iMac models; by comparison, the least expensive model runs \$1,399 and has a 700MHz PowerPC G4, 128MB of RAM, 40GB hard drive, and CD-RW drive. All of them come in the color Apple calls "snow," but the color options will undoubtedly proliferate.

The iMac's 15-inch Apple LCD Active Matrix monitor is attached with a sturdy stainless steel neck. The display accommodates resolutions up to 1,024 x 768, and it's digital all the way, with no RAMDAC involved. You can swivel the display up to 180 degrees and raise it or lower it up to 7 inches.

Although it's rather dated, the iMac's NVIDIA GeForce2 MX video card is a 2X AGP card with an OK chipset in terms of 3-D power. I tested VisionTek's GF2 MX on a 1GHz Pentium III PC awhile back, and it provided a barely acceptable 58.1fps score at 1,024 x 768 but a doable 72.8fps score at 800 x 600.

Apple's primary goal with the current iMac incarnation was to make it a digital hub, and that it did. The iMac comes with five USB 1.1 ports (three on the case, two on the keyboard), two FireWire (IEEE 1394) ports, and a 10/100 Ethernet adapter.

Design. The iMac garnered mixed opinions from folks walking through our lab, but the majority liked (or grew to like) the overall design. The iMac was sturdier than I expected; it weighs just less than 22 pounds, and you can lift it by the neck. The display was extremely easy to manipulate and maintains its viewing angle when moved backward and forward.

You can turn the iMac on its side, flat-panel facing down, and remove four screws from the bottom of the case to see inside. There's little room, precisely enough to add an Airport Wireless networking card or more memory; there's one free DIMM slot, and the iMac can handle up to 1GB of memory. A fan on top and vents around the case's base cool the insides.

Performance. I tried out most of the bundled software, and there was a lot to try. Here are some salient findings: Using iTunes, it took 3:47 (minutes:seconds) to burn 2.5 hours of MP3 files to CD-R and 12:39 to burn a full-length 74-minute audio CD. I used iMovie to make a one-minute slideshow with audio effects. I then used iDVD to tack on a menu and additional effects, and it took iDVD 11:28 to initialize, burn, and finalize a DVD-R. Throughout, I found the OS impressively stable.

The system came with a pair of spherical Apple Pro stereo speakers. Although stylish, their sound was disappointing. Low volume levels were good, medium levels were OK, but the cranked-up volume quality was enough to make an audiophile cry.

Final word. If you're a Mac lover, there's no way you should pass this one up. It comes with a 3-year repair warranty and 90-day (but upgradeable) phone support. ▲

by Cal Clinchard

iMac with SuperDrive


\$1,899

Apple

(800) 692-7753

(408) 996-1010

www.apple.com



COURTESY OF APPLE

Chipsets & Memory Standards



Anand Lal Shimpi has turned a fledgling personal page on GeoCities.com into one of the world's most visited and trusted PC hardware sites. Anand started his site in 1997 at just 14 years old and has since been featured in USA Today, CBS' 48 Hours and Fortune. His site—www.anandtech.com—receives more than 55 million page views and is read by more than 2 million readers per month.

Just one year ago, we could all look to the chipset market for some relief from the non-stop bombardment of marketing over performance from most other sectors of the PC industry. In days when you have to explain to people why a GeForce4 MX is not just a cheaper GeForce4, we sought refuge in the industry where new chipsets were only released when new features and performance improvements could be had. Things are unfortunately changing.

I started to see signs of this when VIA released its recent KT333 chipset for the AMD Athlon platform. In terms of tangible features, the only real benefits the KT333 brought to the table were in the form of DDR333 and Ultra ATA 133 support. The problem with DDR333 support on the Athlon platform is that as long as the Athlon's FSB remains at 133MHz (double-pumped), having a memory subsystem that offers more bandwidth than DDR266 (which runs at 133MHz) won't result in any noticeable performance gains. And for those curious about whether moving to a 166MHz FSB along with a 166MHz (DDR333) memory bus would improve performance on the Athlon XP, it won't. Benchmarks I've run in everyday applications and games show that the performance increase from running an Athlon XP 2000+ at 133MHz x 12.5 with DDR266 SDRAM vs. an Athlon XP 2000+ at 166MHz x 10 with DDR333 SDRAM is between 0 and 6% across the board. This may make sense for the Athlon XP much further down the road as it increases in clock speed, but for now, it doesn't.

SiS has also jumped on the DDR333 bandwagon with their Athlon and Pentium 4 chipsets. The unofficial memory standard (JEDEC has not ratified the DDR333 spec at the time of this writing) does make some sense on the Pentium 4 platform, as the Pentium 4's FSB offers a generous 3.2GBps of bandwidth to the processor. While DDR333 results in a performance improvement over DDR266 on the

Pentium 4 platform, as clock speeds get even higher there will be a need for a higher bandwidth solution. What would actually make the most sense on the Pentium 4 is a dual-channel DDR setup with DDR200 and DDR266 support. A dual-channel DDR200 solution would offer 3.2GBps of memory bandwidth and using DDR266 would kick that up to 4.2GBps: a perfect match for the new 133MHz quad-pumped Pentium 4 FSB.

At CeBIT, NVIDIA joined the list of chipset vendors demonstrating out-of-spec memory support. NVIDIA announced support for DDR333 in their

new nForce 620-D and 615-D chipsets, both of which are dual-channel DDR solutions for the Athlon (Intel hasn't granted NVIDIA a license to use the Pentium 4 bus). The situation with the nForce is a bit different because the integrated graphics can actually use the

added memory bandwidth. NVIDIA did demonstrate DDR-400 at CeBIT, but with DDR333 not even an official spec, you can guess that DDR400 isn't far from overclocked DDR266 SDRAM. NVIDIA even mentioned to us that they didn't see DDR400 as being a realistic option for the next-generation nForce.

In spite of all of these new chipset releases and new types of memory, my recommendations still stick: for the Athlon XP, your best bet is either the VIA KT266A or the KT333 with DDR266 SDRAM. There is absolutely no point to DDR333 right now unless you're running on a platform with integrated graphics that shares your system memory. For the Pentium 4, Intel's 850 chipset still takes my recommendation, as PC800 RDRAM has attained price parity with DDR266 SDRAM. ■

Send your feedback to Anand@cpumag.com

(NOTE: Be sure to check out Anand's chipset and memory comparison chart at www.smartcomputing.com/cpumag/jun02/anandschart.)

There is absolutely no point to DDR333 right now unless you're running on a platform with integrated graphics that shares your system memory.

Sharpening My Teeth

1,600 x 1,200 x 32 with FSAA and Anisotropic filtering turned on is great stuff to showcase the power of my Athlon XP 2100+ avec un GeForce4 Ti 4600 through 3DMark 2001, but what about when I'm not just being a dork and measuring size on my PC? I love to actually put hardware to good use. Games with complex 3-D engines and snazzy effects that showcase 3-D acceleration are a treat. I've recently finished Medal of Honor, Return to Castle Wildenstein, Serious Sam TSE, and Aliens vs. Predator 2 and am working my way through Jedi Knight II: Jedi Outcast. As much fun as I've had, I'm not quite in nirvana and hence I'm going to have to register some BSE-beefs.

First, exactly who do we gamers have to sleep with to get a decent frame rate? Some new games run like that giant tractor that hauls the space shuttle launcher into position! If any newly released game runs with skipping frames on a P4 2.4GHz and a GF4 Ti 4600, then Houston (or somewhere else in Texas), we have a problem.

And I hate to spoil Mr. Game developer's party, but quit messing about with strippers and get some bloody decent AI in, please.

The gamers who fork over the money that pays for those pasties would actually like a bit of a challenge. I'm not asking for brainiac Nazis in Medal of Honor or Kasparov-like planning when Storm Troopers come at me in Jedi Knight II, but wouldn't it be novel for them to do more than just run away when I have anything beefier than a semi-automatic peashooter? I want them to duck, weave, bob, jump, sidestep, and somersault; to have cigarette breaks or pull moonies. Just give me signs of a pulse, and I'd even sacrifice a few polygons in the name of AI. I'm not just talking about first-person shooters here; racing games are guilty, as well. Two player mode is always more fun because your real opponent doesn't drive like the stoned idiot found in almost every racing game that drifts through a pre-scripted racing line. Why can't they drive it like they stole it or cut me off without signaling, like rush hour commuters? Or flip me off at traffic lights? Shall I continue?

Stop with the cheese-cracker story lines, OK? They aren't funny. Furthermore, neither are

bondage-clad zombies that spout key-lime-pie-colored blood. (Yeah, I know you have to make it green to get past the German censors, but isn't there some complicated programming trick to achieve that? <if language=German, then blood=green>.) If you have a good game/movie license, how about trying to stick relatively close to the original movie plot? You know, the thing that made it a success in the first place? Or, do you think you're better than Hollywood scriptwriters? Some games, like JKII, are guilty of overcomplicating the puzzles. Do they all have to be solved with an impossibly difficult Jedi mind trick? And with all the great in-game sounds available, why use the same sexual climax grunt every time your hero picks up some health? What is that about? Collision detection

wasn't invented yesterday, so quit walking characters through walls as if "there is no spoon"! What's wrong with enemies falling to and staying on the floor when they die, like a normal red shirt? How much is it to ask for games to ship relatively bug free? I'm all for hogging bandwidth to download 50+MB patches every so often, but it's

still annoying to have to play a game that crashes, deletes files, or whatever else while I'm playing. Could you make a bit of an effort on characters speaking with accents? Just because you've seen "The Full Monty" 27 times doesn't mean you should do an in-house Brit accent. No I'm not bitter, just looking for some voice acting work. . . . And one more thing: Do you think you could throw some benchmarks into these games? Cheers!

Of course this is just one cranky Brit's opinion. I'm sure my (former game developer) wife will probably make me listen to the Direct X 2.0SDK backwards (again) for this rant, but I'm a long-time game player and reviewer, so I'm going to justify this by saying I'm just standing up for my rights as a consumer and quintessential whiner. ■

E-mail sharky@cpumag.com with your gaming woes, and I'll show you Trespasser.

A black and white portrait of Alex "Sharky" Ross, a man with short, light-colored hair, wearing a dark jacket, looking directly at the camera.

Disrupting Reuters' newswire with a cheery Christmas greeting at age six, Alex "Sharky" Ross became an avid computer user/abuser, eventually founding popular hardware testing/review Web site SharkyExtreme.com. Exposing shoddy manufacturing practices and rubbish-spouting marketing weasels while championing innovative products, illuminating new technology, and pioneering real-world testing methods was just a front for playing with the best toys. The site acquired, he left in 2001. A London native and London School of Economics graduate, Alex currently swims in Silicon Valley.

Who do we gamers have to sleep with to get a decent frame rate?

The Quickening



Kyle Bennett is editor-in-chief of HardOCP.com (hardocp.com), one of the largest and most outspoken PC-enthusiast sites on the Web.

HardOCP.com is geared toward users with a passion for PCs and those who want to get cutting-edge performance from their systems. Beware, though, Kyle is known for his strong opinions and stating them in a no-nonsense manner while delivering some of the most in-depth reviews and PC hardware news on the 'Net.

I was poking around the office storage yesterday and found what was once my fave mainboard, the ABIT IT5H, circa 1997. Nestled on it still is my Pentium 150 that was rock solid at 208MHz. On my way back to my office I opened a package that had arrived that morning, and in it was the new ASUS A7V333 mainboard. I gave the box a good look and realized that the world of computers has changed much in the last five years and is going to change a lot more in the next couple. What I saw on the ASUS box wasn't just a description of a mainboard but a marketing tool for a mainboard with an enormous feature set. Hardware additions include support for 333MHz DDR-SDRAM, ATA133/ATA133 RAID, USB 2.0, FireWire, 6.1 channel audio, and built-in thermal protection circuits to keep you from shaking and baking your CPU. There are even more features to discuss beyond these. If I remember correctly, my ABIT IT5H didn't even have the model number on the box.

The entire product base for PCs is getting better and better all the time and will be migrating to even newer standards soon. We have seen the manufacturers of mainboards, video cards, and peripherals move toward the same goal of producing better products. As they have done this, the basic functions of those products have gotten more and more similar. How are the companies that want your hard-earned jack going to get your attention if their products are like every other out there? Of course, a free-included "busty babes" poster and pirate copy of WinXP was considered by a few board builders, but they soon realized these things could be found easily on the 'Net and those ideas were soon abandoned. Yes, it was decided that we must get more bang for our buck, and we all know more bang is a good thing nine out of 10 times.

There are still purists out there that want only a socket for their CPUs, four DIMM slots, and a South Bridge, as they will argue that all that extra "add-on crap" does nothing but pollute a good enthusiast's or overclocker's board. If you're a purist, get ready to do one of two things, as you are

going to be working with old hardware or giving up your purist hardware religion because the future is almost here.

Many features we see now as "add-ons" are very soon to be "on-chip" implementations, as the chipset makers want that USB 2.0 and ATA133 business.

The ABIT AT7 has been released and is the first "legacy free" mainboard we've seen. Even though it's not truly legacy free, it comes very close. Gone are the PS2, serial, parallel, and gameport connections. What will most likely catch your eye is the absence of a couple of PCI slots, being there are only three on the AT7. Then again, why do we need five PCI slots when we

have on-board support for 12 ATA133 devices, with eight of those supporting RAID, 6-channel sound with optic in and out, a smart card reader, 10/100 Mbps LAN, four USB 1.1, four USB 2.0, three IEEE

1394, and an AGP slot? I only use four PCI slots now in my own box, and three of those are sound, NIC, and extra IDE. Pretty soon, the floppy will be gone, replaced by USB. The IDE will be next, replaced by Serial ATA drives offering tremendous speed increases in a year or two.

High-end video card builders are doing things, as well, even though most are tied to complex reference design cards from NVIDIA and ATI. Companies like VisionTek are beefing up the warranty and support ends for their USA-built products, while companies such as Gainward are shoring up their products with tons of software and extra hardware included, such as FireWire cards, to round out the full multimedia package.

Computing hardware is changing right now and will even more soon. One thing is certain, we will certainly have a few laughs watching those old bastards staring at the backs of their computers trying to figure out where to plug in that Packard Bell keyboard they love so much. ■

Talk with Kyle at kyle@cpumag.com.

How are the companies that want your hard-earned jack going to get your attention if their products are like every other out there?

Introducing the new Palm™ m500 handheld. Inside its sleek little chassis, we've added an expansion slot so you can turn it into the ultimate photo album or eBook. The optional SD cards also let you back up or increase memory, or even access worldwide travel guides. As for included applications, you can download email, import and update Excel spreadsheets, even customize and manage web content with the MyPalm™ portal. We also included mobile connectivity software—add a modem or compatible mobile phone and your information can be accessed wirelessly. It's time to mobilize.



Simply Palm™
palm.com



the last thing you bought
this expandable
had an elastic waist.



Available at: Best Buy Circuit City CompUSA Office Depot OfficeMax Staples

SD expansion cards are sold separately and not included with handheld. SD card shown is an example of available storage capacity. Storage capacity may vary. The m500 handheld requires an Internet account, modem or data-enabled mobile phone and/or third-party software for email and Internet access, sold separately. Screen image is simulated. © 2001 Palm, Inc. All rights reserved. Palm, Simply Palm, MyPalm, Palm Powered, the Palm Powered logo and the Palm logo are trademarks of Palm, Inc. or its subsidiaries. Other products and brand names may be trademarks of their respective owners.

Each month we ask a staff writer to take on our publication editor in a challenge to build the best PC for a certain price. Because our writers don't want to lose their jobs, they always accept this challenge willingly. Tempers will flare. Tools will fly. But only one will prevail.

This month the challenge is to build the **Best Audio System For Less Than \$1,500.**

Samit

I wanted to meet several objectives for this system. I wanted fast, efficient conversion of audio CD tracks to WAV, WMA, and MP3 files (among others) along with the horsepower to easily remix and edit huge files. I wanted a system that would permit quick and error-free audio CDs. The system had to be a solid listening station that could stand its ground without making me long for my home hi-fi setup. As a bonus, I wanted it to run all my other software with ease.

My Insidious mind originally considered skimping on the core system in favor of reference-quality hi-fi speakers and a receiver: a clever ploy, but one I fortunately sidelined. My research made me reconsider several issues, such as shielding, size and space, component and audio quality, and speaker-receiver considerations (properly optimizing speakers with a receiver—trust me when I tell you that burning speaker smell is not fruity). Besides, the more thought I gave to my objectives, the more I realized I was building the best audio PC and NOT the best audio for a PC. We're called *Computer Power User* and not *Stereo Review* for a good reason.


I nearly walked into Marty's speakers at the office. "Pardon me," I said as I veered around the behemoths. I realized that this was Marty "Liberty Lad" Sem's system based on the design I had avoided. Had he used mysterious superpowers to read the Insidious mind? Always curious, I checked for some feedback on the KLH speakers; one reviewer noted sparks and two others noted odors. See for yourself at www.audioreview.com/PRD_119896_1594crx.aspx.

In the interest of buying top-of-the-line PC multimedia components, I had to make some compromises.

For example, I only purchased 256MB DDR PC2700 memory. Any less with WinXP would result in a near-Granny PC (right, Marty?). Also, with the faster processor and RAM, I can rip and write CDs simultaneously (thanks to the BURN-Proof-enabled writer that rips at speeds of up to 40X). Plus I can rip audio from DVDs very easily.

The Klipsch ProMedia 2.1 speakers are universally raved about, but there are no substitutes for a good listen. I paired them with the Sound Blaster Live Audigy Platinum eX (which includes a sexy black breakout box). This package includes connectors for nearly anything you'll ever need (including IEEE 1394).

The CD-RW drive came bundled with Nero 5.5, and Creative included software galore for manipulating audio. The TDK CD-RW audio cable was connected to the CD-SPDIF connector on the Audigy and the Pine 16X DVD-ROM drive was connected to the CD-IN. Given another \$82.84 for my budget, I would have purchased a second hard drive (60GB Maxtor) that I found on sale, but I'm not one to get stuck in a moment I can't get out of. It's a beautiful day and I'm not gonna let it get away. Peace out. ▲



THE PC CHA



Samit G. Choudhuri
Publication Editor
Computer Power User

Component	Model	Price
Case	Antec Performance Series SX1035B with 350W Antec P/S7 *	\$119.19
Motherboard	MSI KT3 Ultra VIA KT333 Chipset Socket A ¹	
Processor	AMD Athlon XP 1600 (OEM) & CoolerMaster Heat Sink & Fan ¹	
Memory	Samsung 256MB DDR PC2700 DDR33 ¹	\$306
Hard Drive	Maxtor DiamondMax Plus D740X—80GB 7200rpm, UltraATA133 (OEM) ²	\$131
Video Card	ATI Radeon 7000 64MB DDR AGP4X (open box return) ²	\$89.60
Sound Card	Creative Sound Blaster Audigy Platinum eX with External Audigy Drive (includes big software bundle) ²	\$208
Network Card	Zonet ZEN3200 10/100 PCI ³	\$3.99
Modem	N/A	N/A
CD/DVD-ROM	TDK 32X/10X/40X veloCD CD-RW (includes Nero 5.5) ^{4*}	\$136.95
	Pine 16X DVD-ROM ^{5*}	\$79.99
Diskette	Panasonic 1.44MB (OEM) ²	\$8
Monitor	MAG 17-inch ^{5*}	\$159.99
Speakers	Klipsch ProMedia 2.1 ⁵	\$179.99
Mouse	Microsoft Wheel Mouse Optical (USB/PS2) ^{5*}	\$19.92
Keyboard	Microsoft Office Keyboard (USB/PS2) (OEM) ²	\$27
Software	MS Windows XP Home Edition (OEM), ² Morpheus, AudioGrabber, MP3Trim, DivX5, Ahead InCD, others	\$88
Other Audio Components	50-pack OfficeMax Spin-X 700MB (32X) CD-R ^{6*}	\$20
Miscellaneous	RoyalComputer.com Testing Fee ¹	\$9
Subtotal		\$1,586.61
Shipping		\$46.53
Tax		\$33.38
Rebates		\$177
TOTAL		\$1,489.52

*Rebate

1 RoyalComputer.com

2 NewEgg.com

3 D.I.T.

4 Buy.com

5 Best Buy

6 Office Max

7 CompUSA

CHALLENGE




6.

Marty Sems

Staff Writer

Computer Power User



Marty

The Rebels could have learned a few tricks from the Empire, such as spend money where it counts (have lots of TIE fighters), but go cheap on the nonessentials (don't give them deflector shields). Here's my latest proton torpedo at Darth Choudhuri's undefeated Death Star.

For the "best" audio PC under \$1,500, I didn't need a fast processor, a good monitor, or even 5.1 surround sound. I did need great speakers, awesome headphones, ample hard drive space, a quiet heatsink fan, Ethernet for broadband downloads, and CD burning/ripping capability. The foundation is there for future home theatre or gaming, but what matters most is high-fidelity sound.

No compact speaker design makes good audio without compromises. It's a matter of physics. Klipsch and Logitech make very impressive use of the tiny boxes most computer users will accept, but for the richest bass, you need sizable transducers with room to resonate. Try KLH 9915 floor-standing speakers with woofers twice the diameter of the Klipsch 5.1's 8-inch driver and tell me which one rumbles your rocks. I added a pair of top-shelf Sennheiser HD 590 Prestige headphones so Judge Blaine Flamig can headbang anytime.

The Hercules Game Theater XP does almost everything a Creative Audigy Platinum does, and danged if I can tell the difference by listening. The Hercules supports digital output and future 5.1 surround speakers to



boot. My favorite toy, though, is the Turtle Beach AudioTron music player. This baby connected to my home network router via CAT5 cable, pulling music effortlessly from my PCs and funneling it to my home stereo. I'm assuming whoever buys an AudioTron already has a network and a real stereo.

Of course, there's no PC Challenge without a few shrieks of hollow despair. After my order languished at WebuildPCs.com for 10 days, I found out the day before deadline that "Ready To Ship" is e-commerce code for "backordered." Heh, heh ROFC (rolling on floor cursing). Adding to the fun, my low-buck 533MHz Celeron crapped out that same night. Huge thanks to Ryan at DIT Computers for the eleventh-hour Duron and mobo swap, and mad props to Michael "Speeding Ticket" Sweet for his audio advice.

Whether Blaine's in the mood for The Replacements, The Plimsouls, or Elmo-palooza, he's going to be blown away. Unfortunately, he's pulling a "Rubin & Ed" on our reclusive butts. Like Crispin Glover's mom in the film, he won't let Samit or me listen to our new tuneboxes until we each go outside and make a new friend. "No friend, no music!" ▲

And The Winner Is...

As I sat at Samit's PC, lullaby bass lines from Bowie's "Ashes To Ashes" swirled around my head, generating visions of that disturbing clown outfit The Thin White Duke donned in the video of said song. It pleasantly haunted my morning. At Marty's rig, I dropped the stellar Sennheiser headphones on my dome and blissfully sank into the springing echoes of The Specials and English Beat (nearly had me doing that one-legged swing-bop thing I had going in 1984). Picking a winner looked impossible.

My head tells me Samit's PC is the better "audio PC." It has the power, memory, and software for fast musical multitasking, and the DVD options give it a definite one up on Marty's machine. We all know about the Klipsch speakers, and the external Audigy was a great touch. Round 1: the Insidious One.

Marty went a bit old school, and my heart loved it. Some hardcore cats around here didn't appreciate his stereo speakers, but I disagree. The behemoths will rattle the neighbors' dishes, which is the whole point. The AudioTron is great, presuming you have a worthy stereo. He pulls back on the wizardry (another drive would be nice), but I dig the volume-stereo-PC approach. Round 2: Marty.

Thus, even though they say it's like kissing a sister, in an unprecedented move, I declare a draw. Ramifications be damned. —Blaine "Wishing Jerry Garcia Was Still Alive" Flamig, content editor

Component	Model	Price
Case	Genica GN200 ¹	\$18.95
Motherboard	Biostar M7VKG (integrated video) ²	\$69
Processor	Duron 1GHz ²	\$54
Memory	128MB PC100 SDRAM ³	\$19
Hard Drive	Samsung SV8004H 5400rpm 80GB ⁴	\$96
Video Card	Integrated S3 Trident Blade	N/A
Sound Card/ Breakout Box	Hercules Game Theater XP ⁵	\$87
Network Card	D-LINK DFE-530TX ⁵	\$9
Modem	N/A	N/A
CD-RW	Samsung SW-224 24X/10X/40X (software and cables added) ⁴	\$80
Diskette	Samsung 1.44MB ²	\$9.95
Monitor	14-inch Acer SVGA (refurbished) ⁶	\$49.99
Speakers/ Tuner	KLH 9915 speakers (\$169.60), KLH Tuner R3100 tuner (\$99.99) ⁷	\$269.59
Mouse	Memorex three-button ⁷	\$7.99
Keyboard	Memorex multimedia keyboard ⁷	\$19.99
Software	WinXP Home full version (OEM), Roxio Easy CD Creator 5 Basic, Winamp, KaZaA, LimeWire, Norton AntiVirus & Ghost 2002, Acid Xpress, MusicMatch Jukebox 5.1, Yamaha XG Studio, Siren Jukebox Xpress, Magix Player Jukebox, others ⁸	\$73
Other Audio Components	Turtle Beach AudioTron (\$273), Sennheiser HD 590 Prestige headphones (\$197.95) ^{5, 9}	\$470.95
Miscellaneous	16-ga. speaker cables, RCA cables, fan/heatsink ^{7, 2}	\$26.96
Subtotal		\$1,361.37
Shipping		\$83.53
Tax		\$42.84
Rebates		\$0.00
TOTAL		\$1,487.74

1 Compgeeks.com 4 Nexthardwareshop.com 7 Best Buy
 2 DIT Computers 5 Newegg.com 8 Amamax.com
 3 1stchoicememory.com 6 Computer Renaissance 9 Sennheiserusa.com

Swappin' Parts

There's A Transformation Taking Place

Each month in "Swappin' Parts," a Computer Power User writer upgrades one out-of-date component in our test machine, MERLE (Mediocre Electronic Refurbished Low-end Equipment). When we're finished we will have transformed MERLE from a silicon trash can into a powerful system we'd be proud to put in our own homes. To date, we've upgraded MERLE's CPU, memory, sound card, speakers, video card, optical drive, case, and PSU (power supply unit).

Many of us pay little mind to our monitors. Why would we? We only stare at them the majority of our waking hours. Case in point: In the introductory "Swappin' Parts" piece we did in *CPU's* December 2001 issue, there's no mention of MERLE's beige 15-inch IBM E54. I'm reasonably certain Michael, who introduced you to MERLE, saw the monitor. I see it myself right there in the accompanying photo to the article, complete with a frown on the screen (presumably to indicate that MERLE's "bad"). I'm more inclined to think the monitor went completely overlooked (just teasing, Michael).

So, I was especially pleased to get an \$800 check to equip MERLE with a monitor of my choice. Considering the market is rife with selection—CRTs the size of Airstreams, slim LCDs, displays that double as television sets, monitors that pivot from portrait to landscape—settling on one (that would satisfy *CPU* editors) would be a challenge. I shopped through hundreds of options, comparing features, performance, prices, and service until I found the perfect package: the Sony CPD-G520 (www.sonymstyle.com). As luck would have it, this was the same model Samit, *CPU's* main man, recently purchased (great minds really

do think alike). And with that, MERLE was prime for a new look.

Why The Sony?

First, the CPD-G520 is flat and big. Specs alerted me immediately to its 21-inch display (19.8 inches viewable). A mere 19 inches from front to back (thanks to a CRT that's 50% shorter than other 21-inch CRTs) means it won't hog desktop space.



Before & After: Turn that frown upside down. The old MERLE (above) gets a face-lift in the form of Sony's 21-inch CPD-G520 (above right).



You can connect two PCs directly to the CPD-G520 or get crazy and add a Mac (it's compatible). A joystick-style front-panel control provides easy access to a comprehensive OSD.

In fact, its footprint is nearly as small as a 17-inch HP seated nearby, saving valuable real estate for yet more accessories and upgrades. Future writers can thank me later.

Of course, there's more to a monitor than size. There's dot pitch, too, and at 0.24mm (0.27mm often being considered acceptable), the CPD-G520 promised eye-pleasing text with precision at a recommended 1,600 x 1,200 resolution and 85Hz refresh rate. Also, Sony's Enhanced Elliptical

Correction System ensured we'd see only quality color display and near-perfect clarity from one corner to the other. Plus, the CPD-G520's MSRP was a reasonable \$799, although we found it for \$645 at Multiwave Direct (www.mwave.com). Dual inputs for hookup to multiple machines (even Macs) were the icing on the cake. For curiosity's sake, we also looked at

Tom's Hardware (www.tomshardware.com) and found only good reviews and recommendations.

If the CPD-G520's \$700 price tag is too stiff, check out the CPD-E540, which at press time was priced at \$599 on Sony's Web site or \$493 at EMS Computing (www.emscomputing.com). You'll have to sacrifice the dual inputs, but you'll still get the good stuff, such as 19.8 inches viewable, 0.24mm dot pitch, and the Enhanced Elliptical Correction System technology.

The Facelift: Going Under The Knife

There's really not much to installing a PnP monitor. I unpacked the box, gathered the various cords, and heaved MERLE's

old IBM E54 into the trash heap out back. (Note to area scavengers: Don't back your cars up just yet. CPU actually donates gently used equipment to local schools. After all, little CPU wannabes have to start somewhere.) I verified that the CPD-G520 and MERLE were turned off and unplugged,

then connected the monitor to the VGA port on MERLE's backside using the bundled video signal cable (for Mac fans there's a Power Mac G3/G4 adapter). I then connected the power cord to the monitor and a power outlet. Power up, and MERLE was ready for action.

At \$1,999, Sony's GDM-FW900 is priced out of our league (but a girl can dream). The FW-900's 16:10 wide aspect ratio is practically perfect and perfectly practical for designers craving side-by-side layout of two full pages or those running multiple apps and windows simultaneously.



Sony's GDM-FW900: One Big, Bad Monitor

Those with a little more coin might consider the Sony GDM-FW900. We did—but who wouldn't with specs like these? When our Sony rep heard about MERLE, she couldn't resist sending us what Sony considers one of its best. Of course, at \$1,999, the GDM-FW900 falls outside of most of our price ranges. But we obliged in order to share some extreme monitor hardware (oh all right, we wanted to check out Serious Sam on a 24-inch landscape, too). Check out the GDM-FW900's goods:

- **Size:** 24-inch wide aspect display with 21% more screen area than 21-inch displays. Panoramic 16:10 aspect ratio provides enough horizontal space to display two pages side by side
- **Max resolution:** Widescreen viewing at 2,034 x 1,440 @ 80Hz
- **Recommended resolution:** 1,920 x 1,200 @ 85Hz
- **Dot pitch:** 0.23mm (minimum)
- **Weight:** 92.6 pounds
- **Dual inputs:** Supports multiple PC usage
- **USB hub:** Connect up to four USB devices directly to the FW900's base

Our installation was clean, though Sony is prepared for such problems as missing drivers and other minor errors with software updates, self-tests, and manuals available from Sony's support section (www.ita.sel.sony.com/support/displays).

Our lab team set the CPD-G520's resolution to 1,024 x 768, which I found painful on a monitor of this size. Icons were huge and text was overbearing. On the other end of the spectrum, the maximum resolution of 2,048 x 1,536 displayed distorted images, hidden taskbars, and stretched text. I settled in at the recommended 1,600 x 1,200, where I found near-perfect text display (even at 7 points) in a Word document, as well as the corner-to-corner clarity that Sony promised. Graphics and animations were also gorgeous with vibrant colors that practically jumped from the screen.

A joystick style control panel provided easy-access to the usual OSD suspects, including contrast/brightness, size/center, geometry, convergence, degauss/moire, color preset and adjust, control locks, language, and auto reset. The OSD is set up frame style, with main menus stacked to the left and submenus to the right. A front panel Picture Effect control lets you choose from Professional, Standard, and Dynamic settings, which incrementally adjust the screen's brightness. Whites appear brighter at Dynamic, dimmer at Professional, and somewhere in between at Standard. I liked how text appeared at Standard, and my current fave, "Moulin Rouge," looked great at Dynamic (though I'm certain Ewan would look good in nearly any lighting).

To CRT Or LCD?

LCDs are sharp, small, and good-looking. So, why go with a CRT? I saw just three unconvincing reasons to consider an LCD. One, the desktop real estate argument: Apparently, Hollywood has had its effect on hardware, too, and even displays can never be too thin. True, the overall LCD surface package looks good and takes up less space, but sadly, LCDs underperform in comparison to CRTs. Behind the LCD's sleek, often silver shell, it just doesn't match the clarity of CRTs.

Two, some say LCDs offer less flicker, but most experts agree that with a 75Hz refresh rate, most users won't notice flicker on a CRT or suffer from so-called subliminal flicker (experiencing eye fatigue and headaches, without actually seeing a flicker). If you do, bump the refresh rate to 85Hz.

Three, energy consumption: LCDs eat less; CRTs eat more. But would you swap your Cheetos-addicted kid for one that nibbles on celery? I think not. Besides, most CRTs go to sleep when not in use, so while you're at work surfing the Web, your monitor at home is happily conserving energy.

A final rant: On top of all that, LCDs only support their native resolutions, plus many require a user to look straight on for a decent view. They're overpriced in comparison to CRTs, and \$645 would have only netted us either a tiny 15-inch screen or a generic 17-inch low-end performer. And although CRTs are heavy (back to the Hollywood thing), the last time I checked, monitors pretty much stay put once unpacked. It's rare to find someone moving one about on a daily basis or juggling them for kids at neighborhood block parties.

Brains & Beauty

MERLE has a new face and it looks beautiful. So far we've given MERLE a new CPU, memory, video card, sound card, speakers, chassis, PSU, DVD-R drive, and, of course, the display. Next month, we bring it all together with a new motherboard. Is MERLE the perfect machine? We're getting there. **CPU**

by Linda Rains

X-ray Vision: 256-bit Encryption

Humans have tried to hide their communications from others for many centuries, and the art of code breaking has been around almost as long as humans have had language. Computers became involved in code breaking during World War II, helping the Allies decipher German codes. Once the computer joined the fray, human-generated codes were no longer secure enough.

With computers involved in almost every aspect of people's lives now, the wars on code breaking have moved to your PC. To protect against crackers, nosy co-workers, and others, sophisticated encryption methods for your computer are available to protect data

in almost any circumstance. And the U.S. government recently decided to greatly strengthen those methods.

Cracks In The Old Code

DES (Data Encryption Standard) is an encryption standard the U.S. government introduced in 1977. Recent experiments have shown the inadequacies of this standard when faced with today's code-cracking hardware and software. To show the weaknesses in the outdated 56-bit DES, organizations recently began building DES Cracker machines. In 1998, for example, the EFF (Electronic Frontier Foundation) used a DES Cracker, built with off-the-shelf components for less than \$250,000, to determine the correct

key in a DES-encrypted file in less than three days. The EFF wasn't trying to steal codes, but it was trying to provide a real-world example of the dangers of not upgrading the encryption standard.

DES Cracker computers, which are fast enough to use a brute force technique to determine the keys, have continued to work more quickly in the last few years, cutting the decryption time for a DES file to several hours or less. (In a brute force technique, the computer tests every possible key combination until it finds the correct key.)

Building A Better Code

To respond to the problem of DES Crackers, the NIST (National Institute

Rijndael At Work

In this simplified model, you gain an idea of how Rijndael creates ciphertext.

Byte Sub Step

1

In **Step 1**, Rijndael reworks each byte of the plaintext, translating it into hexadecimal numbers and creating blocks of bytes.



Shift Row Step


2

Then, in **Step 2**, Rijndael applies the Shift Row process to each block, where the bytes are shifted within their rows.



Billions & Billions Of Possible Keys

Here are the number of possible key combinations under varying strengths of encryptions, including DES, the old U.S. government standard, and AES, the new standard. AES can encrypt data with one of three strengths: 128-bit, 192-bit, or 256-bit. One trillion is equal to 10 to the 12th power, or 10^{12} . You can think of the AES 128-bit key as having 340 trillion trillion trillion key number combinations.



of Standards and Technology) sponsored a competition, asking security and algorithm experts to create a stronger encryption standard. In October 2000, after three years of testing, the NIST recommended that Rijndael (pronounced *rain-doll*) be named the AES (Advanced Encryption Standard). The U.S. Commerce Department officially approved Rijndael as the AES in November 2001, setting the standard for 256-bit encryption. Rijndael was developed by two Belgian scientists: Joan Daemen and Vincent Rijmen.

The AES is a symmetric-key encryption method that can involve 128- or 192-bit encryption. The 256-bit standard will strengthen the privacy of several types of Internet transactions and communications. AES will replace DES for most U.S. government organizations to protect unclassified information. Users can implement AES in software, firmware, hardware, or a combination of

the three. The AES standard was scheduled for full implementation by federal institutions in May 2002.

The Private Sector

Although most transactions on the Internet through the latest Web browsers already use 128- or 256-bit asymmetric encryption, some other security options still use DES for symmetric encryption. The U.S. government expects that the implementation of AES in the private sector, such as for security in items such as PINs from your local bank, will occur steadily in the next few years. Toyota has announced a plan to construct car keys in the next few years that use AES to help prevent theft, for example. (DES followed a similar pattern of steady implementation for everyday use in the private sector after its introduction.)

The U.S. government refuses to make precise predictions on how long AES could remain a viable algorithm. After

all, the rate of improvement in hardware and technology that crackers could use to break AES is impossible to predict. However, the government expects AES to have a life span similar to DES, which was more than 20 years.

Regardless of the future lifespan of AES, you need to deal with it now. You may want to make sure your financial institution and others who are handling your key data have upgraded to AES or an equivalent or have plans to upgrade. With the speed code-crackers are now employing, DES may soon stand for Decrypted Encryption Standard. **COU**

by Kyle Schurman

Wanna Know The Basics?

See www.smartcomputing.com/cpumag/jun02/encryption for the basics of how encryption works, encryption products you can use, and more.

Mix Column Step

3

In **Step 3**, the bytes are shifted within their columns. Each shift in Steps 2 and 3 occurs at a ratio built into the algorithm.



Add Round Key Step

4



=97(hex)=01100001(binary)

mask

=122(hex)=01111010(binary)

XOR RESULT


=27(hex)=00011011(binary)



Finally, in **Step 4**, XOR (exclusive-OR) encryption is applied to bytes in adjoining groups. XOR works by first translating the hexadecimal numbers to binary numbers and then applying a Boolean algebra function (or a mask) to the blocks of bytes. The mask is based on the encryption algorithm. By comparing the digits in the binary numbers of the block and the mask, the encryption program will find "true" or "false" answers to create a new binary number. The program then returns the binary number to its hexadecimal equivalent. Depending on the key length (either 128, 192, or 256 bits), Steps 2, 3, and 4 are usually repeated between nine and 13 times. In the final round only, Step 3 is usually omitted. Also, in the first round only, Step 4 is sometimes applied between Steps 1 and 2.

Got Time?

The U.S. government estimates someone using a DES Cracker computer who could break a 56-bit DES-encrypted file in 1 second...



... would need an average of 149 billion years to crack a 128-bit AES-encrypted file if he were using the same DES Cracker technology.

The Bus Architecture Timeline

Bus architecture has undergone only a handful of changes since Intel built the first microprocessor, the 4004, in 1971.

- 1970** Bus architectures aren't standardized, bouncing between 8 bits and 16 bits.
- 1981** IBM releases its first PC, and the computing industry adopts its 8-bit bus as the ISA (Industry Standard Architecture).
- 1987** IBM introduces the MCA (MicroChannel Architecture) bus to try to supplant ISA. MCA's incompatibility with ISA is a determining factor in the industry's decision to stick with ISA.
- 1988** The computer industry adopts some measures of MCA, including its 32-bit bus, with EISA (Extended ISA). EISA offers data transfer speeds of 33MBps. However, EISA never becomes popular.
- 1992** The VESA (Video Electronics Standards Association) introduces its VL bus (VESA Local bus). It's a 32-bit bus with data transfer rates of 132MBps, but the industry doesn't adopt it as the next generation bus architecture.
- 1993** Instead, the industry goes with PCI (Peripheral Component Interconnect). It's a 32-bit bus that runs at 33MHz and has a 133MBps bandwidth. PCI remains the industry standard today.

3GIO

The Next Generation Of Bus Architecture

When a product has provided you with good service and enjoyment for many years, you hate to let it go. But eventually, someone comes up with a better, faster, and stronger product, leading you to take your Yugo to the auto junkyard or give your first Nintendo gaming console and cartridges to Goodwill.

The same goes for technology that's outlived its usefulness. Your 56Kbps Internet connection served you fine, but you gladly kicked it to the curb when you had the chance to pick up a broadband connection. As you're tossing out the old in your technology closet, you might want to leave room at the curb for your computer's PCI bus architecture. New computer bus architectural designs are coming, led by 3GIO, and PCI is on the way out.

Why A New Architecture?

The PCI-SIG is the standards body that oversees the computer bus architecture, specifically PCI (Peripheral Component Interconnect). Computers have used PCI as the standard for plugging in network cards and graphics cards. PCI also appears in servers.

For about a decade, PCI has been the best way for moving data through, into, and out of a computer. Although the standard operated at more-than-adequate speeds for much of its life span, meeting the speed

needs of today's computer users is becoming increasingly difficult—and expensive—using PCI. PCI's underlying architecture makes it almost impossible to easily increase the speed with which it can deliver data.

The PCI architecture, while it remains reliable, has a few specific problems relative to its age, including a 33MHz bus speed, which is out of sync with other, newer components; a large number of copper wires on the motherboard, which limits placement of other components; and problems with IRQ sharing.

You can think of the various components of the computer, including the PCI bus, as connected by a chain. The speed of the data is determined by the weakest link in that chain, which is currently the PCI bus. And, as you add more and more high-powered devices to your system, such as multiple hard drives or optical drives, the drag of the PCI bus on your overall system becomes more pronounced.

Finally, PCI's plight doesn't figure to improve in the next few years, as components and features such as 10GHz microprocessor clock speeds, high-speed memory, 1Gb and 10Gb Ethernet, and high-speed graphics begin appearing in computers. Intel estimates the bandwidth requirements for input/output could increase by 50 times in the next 10 years over today's input/output demands.

The PCI-X architecture standard is expected to appear sometime during the year, but it's delayed.

PCI-X begins appearing in hardware.

AMD introduces its HyperTransport technology.

1999

January 2001


February 2001

PCI Architecture


VS.

3GIO Architecture

When using PCI bus architecture, several different local buses are in use. A new architecture is needed to meet the ever-increasing input/output bandwidth needs of today's data and applications.



The 3GIO architecture will replace a few architectures currently in use, such as the AGP bus that connects the graphics processor with the memory area. 3GIO also will provide improved speed and bandwidth for several types of needs, including Gigabit Ethernet. The switch is a new piece of hardware included in a 3GIO architecture, and it allows some components to make direct connections to other components through the 3GIO bus.



Therefore, members of the PCI-SIG began exploring a new architecture, and they appear to have settled on 3GIO. Some are calling the 3GIO architecture the most extensive overhaul in the design of the desktop computer since the early 1990s.

Hello, 3GIO

Intel developed the idea of 3GIO in some secrecy through the Arapahoe Work Group. Although Intel did most of the developmental work on 3GIO, Compaq, Dell, IBM, Microsoft, and the PCI-SIG all contributed through the Arapahoe Work Group.

Intel announced the formation of the 3GIO architecture proposal in March 2001. The 3GIO abbreviation came about because Intel simply called the architecture third-generation input/output. Because of the lack of a name initially, some people call the architecture Arapahoe instead of 3GIO. The PCI-SIG has stated that 3GIO will probably not be the final name for the architecture. The group is considering names, but it expects to choose one that will remain consistent with the PCI family of names.

Intel announces its 3GIO standard.

The PCI-SIG announces support for 3GIO.

The Arapahoe Work Group announces the first draft of the 3GIO specification.

March 2001

August 2001

February 2002


Barring complications, 3GIO should replace the PCI standard beginning with computers manufactured in late 2003 or in 2004. 3GIO likely will replace AGP and other internal buses, as well. (AGP, or Accelerated Graphics Port, is the bus that connects the memory area with the graphics processor, providing a fast communication bus. The graphics processor and the memory area previously communicated through the PCI bus, but the AGP provides a faster bus than PCI and alleviates traffic on the sometimes-crowded PCI bus.)

Improvements. The 3GIO architecture will provide some improvements over PCI, including a 75% reduction in the physical presence of copper wires on the motherboard, allowing for more flexibility in motherboard design. External devices can connect directly to the 3GIO bus under the new architecture. And the 3GIO architecture implements a switch, which is a controller that lets devices communicate directly with each other over the 3GIO bus.

3GIO uses different hardware than PCI, but it uses the same software. In this way, 3GIO will provide backward support for the current PCI architectures in use,

allowing PCI software to remain usable without major alterations. 3GIO should be especially important for gamers, who will

see vastly improved performance in their data-intensive software packages. 3GIO also could be a key technology for moving



The 3GIO Architecture

The 3GIO architecture consists of five layers, each of which works together but performs a different function. By splitting the architecture into layers, programmers can make changes to one layer without necessarily affecting the other layers.

Configurations/ OS Layer

3GIO uses the configuration/OS (operation system) layer to properly and efficiently allocate system resources, including memory, input/output space, and interrupts. This layer will optimize the computer's configuration for use with the OS. This layer maintains many of the features and specifications found in the PCI architecture, allowing current OSes to run 3GIO architecture without modification.

Software Layer

This layer ensures compatibility between the 3GIO architecture and various types of software. Using the 3GIO's initialization model, the software layer works with the various hardware components to ensure that the system is optimized for the software. The software layer also ensures that existing software made for the PCI architecture will continue to operate without modification under the 3GIO architecture.

Transaction Layer

3GIO uses the transaction layer in a manner similar to a go-between. The transaction layer receives requests from the software layer for reading and writing. The transaction layer also sends requests to the data link layer to have packets transmitted. This layer can also match response packets from the data link layer with the original requests from the software layer.

Data Link Layer

The data link layer is primarily responsible for reliably delivering data packets across the 3GIO link. This layer must control data integrity as well, by temporarily adding a sequence number and a CRC (Cyclic Redundancy Check) to the packet. If a packet arrives and is marked as corrupted, the data link layer will automatically resend the packet.

Physical Layer

This layer controls the data rates with an initial speed of 2.5Gbps. The physical layer also splits the data between the wires available for data transmission. In the future, speed improvements and new coding techniques will only affect the physical layer, meaning manufacturers can implement the upgrades without affecting the remainder of the 3GIO architecture layers.

Other Players

Other technologies have been mentioned as possible successors to the PCI architecture, but none appears strong enough to match the widespread industry support of 3GIO.

HyperTransport. AMD has sponsored the HyperTransport architecture. Because of AMD's direct competition with Intel in the microprocessor manufacturing industry, some industry experts have thought AMD might try to battle Intel on the architecture standard, fracturing the industry.

However, AMD hasn't developed HyperTransport as a direct replacement for PCI and as a direct competitor for 3GIO. In fact, some reports say the PCI-SIG and other companies tried to encourage AMD to submit HyperTransport as a PCI successor—before 3GIO appeared on the public scene, anyway—but the company declined. As a PCI-SIG member, AMD voted in favor of 3GIO, which probably kills any plans for

HyperTransport to become a PCI replacement.

AMD most likely will incorporate HyperTransport as its data pathway (or front-side bus) for future lines of its microprocessors. HyperTransport will work well for transporting data between internal components.

The HyperTransport Consortium (www.hypertransport.org) will handle future developments of HyperTransport. The Consortium made a late push in July 2001 to become the PCI replacement standard instead of 3GIO, but it didn't work.

RapidIO. The RapidIO standard is somewhat similar to HyperTransport, in that it probably will end up being a high-speed bus architecture that connects other buses inside the computer, rather than being a replacement for PCI. RapidIO will appear mostly in communications and networking equipment, rather than desktop computer architectures.

The RapidIO Trade Association (www.rapidio.org/home) oversees the RapidIO standard.

PCI 3.0. The idea of PCI 3.0 has been discussed for several months, but it probably isn't a competing standard for 3GIO, either. PCI 3.0 has mainly been used to describe the future architecture for PCI, rather than referring to a particular architecture standard.

Actually, 3GIO might eventually become PCI 3.0, depending on how the PCI-SIG decides to name the final standard for the architecture.

PCI-X. PCI-X is the fastest version of PCI currently available. However, the 3GIO standard will be much faster than PCI-X. Individual 3GIO wires will carry about 12 times more data per second than individual PCI-X wires initially. Future upgrades to 3GIO wires will increase the difference in data transfer rates even more. Another advantage 3GIO has is that 3GIO wires are less expensive to use and build than are PCI-X wires.

PCI-X meets the needs of current computer setups, but it has nearly reached its data transfer rate limitations. The availability of 3GIO will hasten the end of use for the PCI-X

standard. If the experts are correct, and 3GIO hardware begins appearing in 2003, PCI-X will have had a short life span of only about 24 months. PCI-X initially was expected to appear sometime in 1999, but delays cost the architecture nearly two years before its introduction.

Those involved with PCI-X say they didn't introduce it to the PCI-SIG quickly enough, letting it gain approval from all companies involved in the PCI-SIG. Those working with 3GIO have learned from the problems and mistakes made with PCI-X, which explains the speed with which the PCI-SIG approved 3GIO.

PCI-X 2.0. This standard is aimed at server hardware, but 3GIO will probably replace it, too. PCI-X 2.0, also called PCI-X DDR and PCI-X 266, might hold on a little longer than PCI-X because server manufacturers typically are slow to adopt new technologies. The PCI-SIG also governs PCI-X 2.0. Like PCI-X, PCI-X 2.0's overall data transfer speed is plagued by its parallel connections instead of the serial connections 3GIO uses. ▲

the large amounts of data for video and images quickly and efficiently through the computer's graphics systems.

The PCI-SIG steering committee voted in August 2001 to begin the process of replacing PCI with 3GIO. Because the steering committee consists of representatives from industry giants such as AMD, Compaq, Hewlett-Packard, IBM, Intel, Microsoft, Phoenix Technologies, and Texas Instruments, widespread adoption of its recommendations seems like a foregone conclusion. In February 2002, the Arapahoe Work Group released the first draft for the 3GIO specification, which the PCI-SIG was reviewing at the time of this writing.

3GIO's Role

Nearly all experts think 3GIO will be a great architecture for delivering data to and from network and graphics cards. However, some questions exist about how 3GIO will fit into the overall architecture of the computer. USB and FireWire already handle connections between the computer and many types of peripherals, including digital cameras and printers. Because implementing different architectures within the computer can be an expensive process, some experts think migrating the PCI architecture toward the USB and FireWire standards would be a better solution.


It will be interesting to see how 3GIO's implementation progresses over

the next several months. Although 3GIO appears to have support from plenty of industry giants, the politics of developing an industry standard are sometimes mystifying.

Still, 3GIO appears to provide the best hope for improving on the steady but aging PCI architecture. As much as some PCI fans don't want to admit it, it's time to retire the PCI bus in favor of an improved architecture. **CPU**

by Kyle Schurman

See the "Buses & Chipsets" article on page 64 for more information about buses.



**OUR 270-HORSEPOWER ENGINE CAN BEAT UP YOUR...WAIT,
YOU DON'T HAVE A 270-HORSEPOWER ENGINE.**



One-year OnStar Safety and Security service contract is standard on TrailBlazer LT and LTZ. Call 1-800-ONSTAR-7 for system limitations and details. *Based on horsepower and GM Medium Utility. Based on GM testing. Official EPA estimate not yet available. OnStar is a registered service mark of OnStar Corporation. ©2001 GM Corp. Buckle up, America! For more information, visit chevy.com

Ouch. The truth hurts. But it just so happens the 2002 Chevy™ TrailBlazer® is the world's most powerful midsize SUV.* The driving force behind all this muscle? The all-new Vortec™ 4200 six-cylinder engine that serves up an exhilarating 270-hp, more standard horsepower than *any* V8 competitor. It also

INTRODUCING THE ALL-NEW CHEVY TRAILBLAZER

dishes out a whopping 275 lb.-ft. of standard torque. Which is more than (hey, big surprise) *any* six-cylinder competitor. On top of that, you'll be able to tow up to 6,400 lbs.† worth of boat, trailer or anything else with or without an engine.

You'll get six-cylinder fuel efficiency.‡ And you'll get smooth, quiet, refined engine performance, thanks to the Vortec 4200's inherently balanced in-line design. (See, extraordinary power isn't *all* about brute force.) Test drive a new Chevy TrailBlazer. After it, you'll find that everything else just seems kind of weak.

TRAILBLAZER  **LIKE A ROCK**

segment. Excludes other GM vehicles. †Maximum trailer ratings are calculated assuming a properly equipped base vehicle plus driver. See The Chevy Trailering Guide for details. **Estimated MPG 16 city, 23 highway for 2WD models. or call 1-866-TRAILBLAZER.

Evolution Or Revolution?

Your PC In 2003

There are no crystal balls in the computer industry. Recessions often sweep new technologies away and manufacturers are usually hesitant to give the competition even a hint at their plans. Nevertheless, we spoke with a number of leading companies to piece together what a PC in 2003 might look like. In the next 12 to 18 months, expect more evolution than revolution—with some surprises.

Hard Drive

Darrin Bulik, Western Digital's (www.wdc.com) technical marketing manager, says IDE hard drives in 2003 will still be 7,200rpm with three-platter designs, but each platter will feature up to 80GB capacities. Drives will have an 8MB or higher cache option. Also, Serial ATA will take hold, offering 150MBps burst rates and sustained transfer rates "into the high 60s [MBps] at the outer edge of the platter."

Eckert and Mauchly create the UNIVAC, the world's first for-sale computer system. It uses 5,000 vacuum tubes and weighs a svelte 8 tons. The U.S. Census Bureau buys the first unit.

IBM launches the System/360, the first mainframe system to use interchangeable hardware and software.

The first commercial PC, the MITS Altair 8800, is released, utilizing an Intel 8080 processor and a version of BASIC written by Bill Gates and Paul Allen.

Steve Jobs and Steve Wozniak start Apple Computer.

1951

1964

1975

1976

by William Van Winkle
Graphics & Design by
Erin Rodriguez & Ginger Riley

CPU

Speeds will easily exceed 3GHz. Intel (www.intel.com) should release its "Prescott" architecture in late 2003. AMD (www.amd.com) will unveil its Hammer line. (For more on CPUs, see page 58.)

Removable Storage

Tape may make a comeback, as full-drive backups become increasingly impractical with optical burners. Brad Renfree, director, product line management, Seagate Removable Storage Solutions (www.seagate.com), says "We have extended the existing Travan technology, moving from 20GB compressed to 40GB compressed. This is really significant because it's been a while since tape has kept up with disk technology, and 40GB is now the sweet spot in hard disks."

Motherboard

Samuel Kang, EpoX (www.epox.com) spokesperson says, "By 2003, most boards will not have the PS/2, serial, or LPT ports; they will all be replaced by USB ports. USB 2.0 is here to stay. You'll also see integrated 802.11b, 802.11a, and Bluetooth, as well as 128-bit architecture." He adds that ATA/166 will come in Q4 2002 and integrated IDE RAID will replace SCSI.

Chassis

Intel is spearheading a new "Tidewater" design format for low-end PCs that will shrink the micro ATX design even further. At the high end, expect few changes in the next two years. However, as the amount of heat that systems generate continues to increase, look for liquid cooling to be the "next evolutionary step in chassis enhancement," says Tim Hunting of Koolance (www.koolance.com).

Peripherals

John Skeeahan, Logitech (www.logitech.com) senior global marketing manager, says, "expect to see innovative Bluetooth devices appearing across the peripheral landscape..." An early example is the Cordless Presenter, Logitech's first Bluetooth device, which doubles as a wireless, programmable mouse and a laser pointer.

RAM

"Currently there is a rapid shift from SDRAM to DDR as the dominant memory technology," says Jaja Lin of Kingston (www.kingston.com). "PC133 demand... will drop off sharply in 2003. PC2100 DDR will dominate through mid-2003. By the second half of 2003, PC2700 DDR is expected to replace PC2100 as the mainstream memory technology."

Monitor

Usability will be the big shift in monitors, and LCDs will grow in popularity as prices drop. Todd Fender, NEC-Mitsubishi (www.necmitsubishi.com) LCD product manager, says monitors will become more intelligent, "doing things for the consumer automatically." At the low end, "we'll see DDC/CI (Display Data Channel/Command Interface), the next generation of PnP for displays. This involves two-way communication... to control all of the settings previously handled by the front buttons." He says new designs will feature a 13mm bezel frame and autobright functions.

Optical

Mark Becknauld, product manager for Yamaha's Multimedia Products Division, says CD-RW and recordable DVD drives will increase in speed and popularity, plus "look for high-performance USB 2.0 and IEEE 1394 external burners... to gain in the market. With an external, you only need to buy one drive for all occasions rather than a burner for every system in your possession."

Audio Adapter

External connectivity will let people upgrade their audio who wouldn't have done so before, says Greg Thorn, director of audio marketing for Creative Labs (www.soundblaster.com). He adds, "thanks to USB 2.0 and IEEE 1394, the external pipe is only going to grow broader, and that will enable us to provide even more value."

Video Adapter

ATI (www.ati.com) senior product manager Toshi Okumura says ATI will launch integrated mobo chips and discrete add-on card products in quick product cycles. Mainstream chips for performance users are planned six months later. "We probably won't have true realism in 2003, but we'll be closer. You will see video encoding migrated into graphics chips, just like DVD decoding did before it." Okumura says a Pentium 4 platform product code-named "RS200" is slated to succeed the A3 chip for 2003, and the "R300" will supplant the Radeon 8500. (For more on video adapters, see page 64.)

Speakers

You may have to wait beyond 2003 for digital speakers. Mark Casavant of Klipsch (www.klipsch.com) says, "A lot of manufacturers are pursuing digital speakers for digital's theoretical clarity. But to hit their low price points, these vendors are trading component quality for the integrated digital-to-analog circuitry... we've focused on keeping component quality as high as possible in our analog offerings, then using an external digital decoder, our DD-5.1, to complete the DAC conversion," which should be a "popular design approach in the performance market."

Chipset

VIA's (www.via.com.tw) Richard Brown says expect FSB speeds up to possibly 800MHz in 2003 and "support for DDR-II-SDRAM initially at 400MHz, higher bandwidth interconnect with the South Bridge (our V-Link will be up to 1GBps), and an AGP 8X graphics port."

He adds that the real excitement will be in integrated 802.11b support (via an external receiver), Serial ATA, and an integrated USB 2.0 controller. (For more on buses and chipsets, see page 68.)

The Xerox 8010 Star system debuts. Xerox later invites Apple engineers to examine the new design, allowing Apple to "borrow" the concept and pioneer a popular adoption of the 8010's GUI.

IBM announces the IBM PC. Other manufacturers "clone" its design, assuring the PC platform's success but ultimately dooming IBM's efforts to rule personal computing.

Apple releases the GUI-based Macintosh, ushering in a sweeping change in how users communicate with personal computers.

Microsoft releases Windows 1.0, but it's not until Windows 3.0's release in 1990 that Microsoft gains any traction in the GUI space.

1981

1984

1985

Today's PC

The future is meaningless if we can't appreciate the present. Here's a look at what today's PCs offer.

- **CPU:** 2GHz Pentium 4 and Athlon XP chips are becoming mainstream.
- **Motherboard:** Whether it's Intel's Socket 478 for the Pentium 4 or AMD's Socket A for the Athlon family, today's mobos increasingly feature integrated audio, modem, and Ethernet capabilities. Legacy ports are still standard, USB 2.0 is emerging, and Integrated RAID is a high-end niche feature.

- **Memory:** 256MB to 512MB of DDR-SDRAM is common, as RDRAM edges closer to obscurity.
- **Chassis:** ATX midtowers with an adequate power supply and one or two extra case fans remain the staple.
- **Video adapter:** NVIDIA's GeForce3 and GeForce4 MX, as well as ATI's RADEON 7500 and 8500 products, dominate mainstream AGP 4X graphics chipsets.
- **Monitor:** Sub-\$200 17-inch CRTs are the norm with new systems, while 15-inch LCDs dipping under \$300 may herald a shift to flat panels in the mainstream.

- **Storage:** The average ATA/100 or ATA/133 hard drive holds 60GB and sports 7,200rpm.
- **Sound adapter:** Creative Labs' SoundBlaster Audigy pushes the company's Live! family down from the performance to mainstream market segment. Surround sound and 3-D spatial effects are common for the entry-level market.
- **Speakers:** Respectable 2.1 sets sell for less than \$50, with high-end 5.1 systems reaching \$400. Analog speakers remain the norm; flat-panel speakers are little more than a curiosity.
- **Optical:** CD drives are gradually giving way to 24X or 32X CD-RW drives. Combination

CD-RW/DVD drives are becoming increasingly common. Recordable DVD drives are well under \$500 (street), but format concerns still keep most consumers away.

- **Network:** 10/100 Ethernet reigns as Gigabit Ethernet taps away at the corporate world. 802.11b is the wireless standard of choice, and prospects seem bleak for phoneline and powerline products. ▲

(For a look at how today's hardware is being implemented into the home, check out "The Home Of Tomorrow" at www.smartcomputing.com/cpu/mag/jun02/home)

The Changing Faces Of OSes

Getting official statements about future OSes is like pulling teeth, but here's an idea of what new releases may offer.

Apple's Bill Evans says, "open standards and platforms, that's the direction that OSes have needed to go . . . We are basically the platform of choice for Java. Look at the environment that Max OS X is creating and extrapolate from there." He adds, "We're working hard to

foster a platform where two people in a garage can make the next killer app."

The Aberdeen Group recently stated, "The development of Linux is moving faster than any other commercial operating system." (It also has the most ground to make up on the desktop.) Bill Mason, Red Hat's director of product management, says, "During 2003, look for enhancements to Red Hat

Linux in the areas of graphical installation capabilities and configuration tools," plus improvements to the desktop display.

Microsoft is mum about future OSes, but Paul Thurrott's SuperSite for Windows (www.winsupersite.com) mentions "Longhorn" is due in late 2003 or early 2004. "Blackcomb" is expected in 2004/2005 with integrated .NET functionality.

The big push now may be extending Windows throughout

the entire house. "[By] harnessing the remote Desktop and wireless networking features of Windows XP and Windows CE .NET, 'Mira'-enabled smart displays will deliver to consumers the freedom of the complete experiences in Windows, including browsing the Web, sending or receiving e-mail messages, listening to music, and editing and displaying digital images, from any room in the home," reads one Microsoft statement. ▲

Linus Torvalds, a Finnish grad student, posts to a newsgroup that he's working on an open source kernel as a "hobby." He publicly releases the source code for Linux that year.

Microsoft launches Windows 95. Major innovations include 32-bit processing, Internet access, and efficient multitasking.

IBM's Deep Blue beats chess master Garry Kasparov in two matches (Kasparov eventually wins the series).

Microsoft releases Windows XP. Viewed by many as either a cosmetic upgrade to Windows 2000 or a significant evolution of the Windows product line.

1991

1995

1996

2001

Affordable. Recordable. And Very, Very Portable.

Introducing the Memorex Pocket CD-R™. With 185 MB, it holds hours of MP3 music, 140 high-res photos or tons of data. At about three inches, it fits any pocket. At pennies a meg, it fits any wallet. With both writeable and rewriteable formats, it's big enough to do almost any job, yet small enough to go almost anywhere. For a closer look, go to memorex.com. Because this little CD will be one of the biggest things in recordable media.



Memorex®
Is it live or is it Memorex?™

©2002 Memorex Products, Inc.



Processors

The Next Little Thing

You can safely assume that your CPU will experience significant changes in the coming decade. As engineers overcome power, heat, and size obstacles, new technology will arrive on your chip. One of the people wielding significant influence over those innovations is Intel CTO, Pat Gelsinger, who gave us this prediction: "By 2010: 30GHz, 10 billion transistors, and a tera-instruction per second." (See our interview with him on pages 62 and 63.)

Now, Intel is not known for timidity when it comes to long-term projections. In fact, the company is over-betting the SIA's (Semiconductor Industry Association) rather conservative 2010 technology roadmap by about 18GHz, but Intel does tend to keep mum about its immediate future roadmaps. AMD, however, leans pretty far the other way, publishing its CPU roadmap right on its Web site. But the company is very close-mouthed about its long-range (currently anything post-eighth generation) technology plans.

BBUL: Check Your Packaging

The size of your package is no laughing matter. It is an essential, integral structure that feeds power to the chip and draws heat away, transfers information into and out of the mainboard, and protects the chip from possible environmental

damage. In Intel's case, its goal of reaching a 30GHz microprocessor sporting 1 billion transistors isn't going to happen without technological innovation in package design. The company is banking on a technology in development called BBUL (bumpless build-up layer). The breakthrough could come with packages that are built up around the silicon die as opposed to manufactured separately and then bonded together.

BBUL isn't the most exciting item on Intel's list, even according to Gelsinger, but it will have an impact on many other innovations currently keeping design engineers awake at night: "I'd call BBUL sort of a meat and potatoes kind of technology; it just makes everything better. It's going to enable greater integration, improved characteristics, and new combinations of technologies coming together, and it's going to be really core to the product line over time."

Package primer. The Pentium 4 and its 55 million transistors is packaged in a flip-chip BGA (Ball Grid Array) package connected to the silicon die by a set of C4 solder bumps. These bumps are the last structures put on and are used to connect the silicon to the package.

The package makes an electrical connection between the pins and bumps.

Manufacturing of the package begins with the core (middle layer), consisting of a sheet of plastic with copper interconnects, called vias, passing vertically through the core. Holes are laser-drilled and filled with copper so the vias (as many as 10,000 in today's CPUs) can be housed in the core. Copper interconnect layers are then assembled on either side of the core ensuring that each C4 bump is routed to the correct pin. Finally, when all pins are attached, the package is set for insertion into the board. In the P4's BGA case, the pins are replaced by balls that attach to the surface of the board. The limitations of this technology will become obvious as microprocessors increase in complexity requiring that ever more C4 bumps be placed on the shrinking package.

BBUL technology enables higher performing and/or lower power consuming chips to be manufactured. It removes the need for a top interconnect layer by embedding the silicon directly in the package core, rendering the C4 bumps obsolete and making for thinner and lighter packaging. Only a lower interconnect layer is required to be built up. The lack of C4 bumps allows closer connection and improved route signals across different parts of the die. Shorter interconnects reduce inductance. More efficient power performance is possible by capacitors situated on the pin side of the package that are closer to the silicon, allowing it to run at higher speeds. Chips operating at lower voltages will be possible through diminished electrical noise.

Intel's 8008 has an 8-bit bus, 3,500 transistors on 10-microns, and a speed of 800KHz.


Intel introduces the 4004. It has a 4-bit bus and is the first single-chip CPU. The 4004 has 2,300 transistors and runs at 108 to 740KHz.

The 8-bit Intel 8080 serves as the brains for the Altair \$395 PC Kit.

The 8-bit version of the Intel 8088 powers the IBM PC and puts Intel on the map with 5MHz!

AMD introduces the 8080A, a reverse-engineered version of Intel's chip, and gives away a television at the fifth anniversary employee street fair.

1971 1972 1974 1975 1979



Intel is hoping its BBUL (bumpless build-up layer) technology, shown here in comparison to the current FCPGA (flip-chip pin grid array) packaging, will help the company reach its 30GHz goal.

Reducing power dissipation in tandem with voltage is of prime importance in the quest for eking more MHz from the finite limits of physics.

Intel is aiming to physically shrink the CPU down in size to little more than the width of a dime with smaller and shorter pins. This will be necessary in the future for smaller form factors such as mini-mobile PCs, cell phones, and PDAs, and most especially when we are all wearing Star Trek communicator pins (more about that in the "Silicon-Based Radio" section). Even more intriguing is that BBUL could allow for multiple chips in the same package. An interconnect layer can make connections from die to pins as well as among the die. Embedding a CPU and chipset into one package is another way to decrease the footprint. Ultimately, the idea is to lay out the CPU, chipset, GPU, memory, and a few other amazing goodies onto a single 1mm-thick substrate.

AMD's resources are not quite as hefty as Intel's 900-strong IC-packaging workforce. But having recently released its first organically based substrate CPU with the Athlon XP, AMD has shifted away from the limitations of ceramic packages. With 0.13-micron and SOI (Silicon-On-Insulator), AMD's eighth-generation Hammer processors will be able to scale better; AMD is playing the packaging game, too. Companies such as NEC and GE have fired blanks in bumpless-bonding chip packaging for military applications, so the nut has yet to be cracked.

Intel still faces some challenges to attain acceptable yields, making the entire process cost effective for mass production. Then, and only then, toward the second half of this decade, will we likely see BBUL-based CPUs in pocket PC-sized devices that transfer data to and from your solar-powered home, your hybrid convertible car, and your Intel-stellar office. Don't

worry, though; if desktop PCs still exist, the prospect of playing Quake 9 on a 30GHz BBUL-based CPU will have its own reward. So, how far off is all this? According to Gelsinger, "It's medium term, several years away yet, but I characterize that as on track. The timeline has been laid out . . . and we're still pretty comfortable that the technology is maturing at that rate. Where others have failed with techniques like this, our confidence is building that we truly are going to be able to deliver what was described last year."

Viva la Transistance!

Advances in transistor technology are key, as those little blighters are the silicon-based circuits that process the 1s and 0s on CPUs. The crème de la crème of transistors are currently up at the 1 billion operations per second mark, but the TeraHertz transistor will be much nippier and capable of switching between its "off" and "on" state 1 trillion times per second. But as more high-speed transistors get slapped on to a sliver of silicon, even the most elaborate of cooling techniques won't be enough to cool things down, let alone deal with high-power consumption. (Get in-depth TeraHertz transistor information on pages 44 to 47 of the March 2002 issue of *CPU*.)

At a recent demonstration of the TeraHertz transistor, Gerald Marcyk, Intel's director of components research, said, "Smaller and faster just aren't good enough any more. Power and heat are the biggest issues for this decade. What we are

The Intel 80286 is introduced with 134,000 transistors running at 12MHz. Intel also releases the 80186 but not for commercial use.

Transistors reach 275,000 in the 1.5-micron Intel 386, which runs at 16MHz. It has a 32-bit bus and multi-tasking capabilities.

The 25MHz 1-micron Intel 486 has 1.2 million transistors. It's the first CPU with a math co-processor.

The AM386 from AMD has 200,000 transistors.

Intel's 66MHz486DX2 with 1.2 million transistors is introduced.

The AMD AM486 reaches 1 million transistors and speeds of 33MHz. The 0.8-micron Intel Pentium has 3.1 million transistors introduced at 66MHz.

1982 1985 1989 1991 1992 1993

doing with our new transistor structure is helping make devices that are extremely power efficient, concentrating electrical current where it's needed."


Although Intel already possesses chemical vapor deposition tools that can deal with the high k material for use with

limitations associated with it. . . . Though Intel's use of a fully depleted substrate transistor with a "high k gate dielectric" has been demonstrated at 2.6THz, AMD actually has the fastest (currently) .015-micron transistor at 3.3THz. IBM's 2THz silicon germanium

they posed the question: Why not use CMOS for the radio?

A smart radio circuit with an automatically reconfigurable wireless network connection is really what they had in mind. In fact, the initial vision was to walk out of Starbucks during an 802.11 session, have the device hand off seamlessly to a cellular network, and ultimately incorporate automatic "sniffing" capability so the device always knows what networks are available and is ready to connect to them automatically. All this within a highly stringent 2X to 3X criteria, meaning that compared to current solutions (think 802.11), the device can't consume more than two to three times the power or cost more than two to three times the price. This is the vision. The reality so far breaks down into two major design problems: the radio transceiver itself, and the controller that enables those nifty multiple protocol connections. But acceptable sniffing performance won't be a stroll in the minimall, either.

In the lab. Steve Pawlowski, director of Communications & Interconnects Technology, one of the engineers mentioned above, and the Intel Fellow who authored the CMOS Radio white paper, was kind enough to give us this status report from inside the research lab where the development platform is emerging. Using high-powered processors, DSP (digital signal processing) techniques, and high-speed digital logic, engineers are constructing a system to be reconfigurable for (currently) 802.11a and an experimental UWB (ultra wide band) prototype application operating in excess of 100Mbps.



This size progression of transistors shows a continuation of the shrinking trend.

300mm wafers, others in the industry are currently moving in the direction of partially depleted SOI. Gelsinger says, "We talked a lot about SOI, but the approach today is a partially depleted substrate, which gets you some benefit but at substantial cost, and the implication of that for mass market approaches is that it's going to be later and more expensive. . . . In contrast, what we have done with the TeraHertz transistor is a locally fully depleted substrate that gives you significant scaling and much greater benefits. So we think that the interim technologies, while there are going to be applications for them, I don't want to imply that there won't be some places where they make sense, I view it more as a sort of a niche technology, but the mainstream isn't going to go that way just because of the cost and manufacturing

(SiGe) technology ("double gate" transistor) is also in the works.

Silicon-Based Radio

Although packaging and transistor cramming are important technologies, a new addition to your future CPU might have the biggest impact on how you compute. It started as a "skunkworks" project involving an Internet tablet concept device containing a display, a wireless connection, and a little bit of computing to tie it all together. What became the CMOS Radio project was initially driven by a couple of engineers who wanted to build on the initial product concept and make it more robust. The problems they saw involved high cost, limited range, and inability to move between different types of wireless networks. Somewhere along the way,

100MHz is cracked when the Intel Pentium on 0.6-micron arrives.


5.5 million transistors appear in the 0.6-micron Pentium Pro, the highest speed to launch at 200MHz.

AMD introduces the AMD K5. It has 4.3 million transistors and a top speed of 75MHz.

Intel introduces MMX with the Pentium 200MMX, the Pentium II Klamath (0.35micron), and Deschutes (0.25micron), which cracks 300MHz. AMD's 0.25micron K6 is released at 233MHz

AMD introduces the K6-2 with 9.3 million transistors. Intel releases the 0.25-micron Celeron 300A (Mendocino), which has 19 million transistors and becomes an overclocker's favorite.

1994 1995 1996 1997 1998



Shown here is a proposed hybrid architecture using the MCA (microcoded accelerator).

devices, and circuits and apply them to their particular product. The circuits that we build to validate and prototype our designs and technology may not be the circuits that they end up using. So they have access to the circuits we've built, but what we're really developing is the capability and the expertise of the process so they have a library of tools available to go do their product development."

MCA. The truly tricky bit is the reconfigurable MCA (microcoded accelerator) that handles the computations necessary to cope with multiple existing protocols and new ones developed in the future. "Instead of one huge pipeline that tries to do all the computation, we broke down the algorithms into constituent components and then determined the types of computation each one of those components would need. For example, a Viterbi decoder needs different types of computation than a FIR filter. So we broke the protocols up into their constituent elements and we are constructing a pipeline and an ALU (arithmetic-logic unit) for those particular elements. The plan is to make it configurable, to use the same hardware and refigure the "control logic" to be able to control that hardware in a different manner." So these constituent elements are connected together in a "hybrid MESH network topology that can be programmed to execute different communications solutions at an energy-per-computation cost less than that of current DSP solutions." Voila! The result is a very flexible math machine

This is where marketing goes out the window and the difficult and exhilarating work of technical innovation comes alive: "We are trying new things to compensate for the effects of CMOS as it scales in terms of process and lower voltages so that we can eventually integrate the entire radio subsystem, including power amplifier, low noise amplifier, A/D converters, etc. We want to take the entire radio front end and put it in a piece of silicon in order to get the cost down. So there are some interesting challenges. With some of the other technologies, you can use 2.5- to 3-volt power supplies so you get to take advantage of higher voltage circuits (i.e. a high efficiency Power Amplifier) and better immunity to noise that you don't get with a 1.2- or

1-volt process, if it ever gets to that point. But it's an engineering problem, not really a physics problem that can't be solved." Steve has been bopping around Intel for the last 20 years and has accumulated more than 30 patents along with the hard-earned opportunity to play with such challenging toys, especially if they have the potential to become ubiquitous, play on Intel's core competencies, and can continue to pass muster through regular reviews.

"From the CMOS side, we are in the innovation stage in terms of the circuit development. We have several basic designs and test chips in process. We're characterizing those designs and building the appropriate CAD models so the product group could use those designs,

AMD cracks the 1GHz barrier first with a 0.18-micron K75 core sporting 22 million transistors. Intel strikes back with Pentium III 1GHz. Both companies shift from slot to socket based CPUs. Intel recalls the

Pentium III 1.13GHz. 0.18-micron Pentium 4 (Willamette) launches at 1.5GHz with 42 million transistors. The 0.18-micron Athlon (Thunderbird), with 37 million transistors, also approaches the 1.5GHz mark. The AMD value-aimed Duron challenges the Celeron.

The AMD K6-3 is released at 0.25. The Intel Pentium III with 9.5 million transistors hits 500MHz and then moves to a 733MHz 0.18-micron Coppermine core with 28 million transistors. AMD's 0.25-micron K7 architecture squeezes the performance delta between the two companies for the first time.

AMD introduces the 0.18-micron Athlon XP (Palomino) with 37.5 million transistors. It lasts into 2002 at 1.73GHz.

The Duron breaks 1GHz. Intel releases the first 0.13-micron processors with the Pentium III (Tualatin) and Celeron (Tualatin). Intel breaks 1GHz with Pentium III on the mobile side. The Pentium 4 breaks the 2GHz mark and moves from FC-PGA to micro-BGA package.

Intel shrinks the P4's die to 0.13-micron with the Northwood (55 million transistors). It reaches speeds of 2.4GHz and moves to a 533MHz FSB. AMD shrinks the Athlon XP die to 0.13-micron with the Thoroughbred core and demonstrates the ClawHammer.

1999

2000

2001

2002

capable of conversing in multiple protocols of RF communication that powers up only the necessary areas as needed.

"The MCA work is still at what we call the pathfinding stage. We're just getting started in terms of looking at the architecture, how we would interconnect these things, the software, the runtime environment, the development tools, everything. We've been doing work for a while, but it's a longer term project and still early in the discovery stage."

As Gelsinger says, this won't all happen at once: "I very much see incremental changes along the way. CMOS Radio, for example. Today you have cell phones with 60 to 70 odd components in there and three big ICs [integrated circuits]. I definitely see that there is going to be an incremental step . . . where we sweep most if not all that stuff together into the analog chip. That will be one integration, but that's not all the way to the 'Radio Free Intel' vision of having all that on just a corner of a die."

Although BBUL, TeraHertz transistors, and CMOS radio will all have an impact on the structure and functionality of future CPUs, innovations in photonics, sensor networks, and MEMs (microelectromechanical systems) are also moving toward silicon. What is now considered a CPU may become only a tiny subset of

what is on the processor of 2010, and engineers in labs and universities and junior high schools all over the world are just itching to show us the rest.

On The Road Again

30GHz and CMOS Radio can't come soon enough? We've trawled the usual sources to bring you the skinny on this year's anticipated desktop CPU harvest. Intel's Pentium 4 (Northwood) 2.4GHz is the current performance champ. Better performing P4s, with 533MHz FSB speeds (133MHz quadpumped) with the added bonus possibility of PC1066 RDRAM support are set to go live around the time you read this. Intel has already publicly demonstrated a P4 at 3 and 4GHz, so there's life in the ol' P4 yet.

But before gamers jump ship from an AMD Athlon XP to a P4, those 300mm thin wafers packing 0.13-micron Northwoods need some drastic price cuts. Also in Q2, Intel is expected to launch a range of new desktop chipsets supporting 533MHz FSB, integrated graphics, and integrated USB 2.0. For value-minded folks, a new Celeron line with SSE2, the P4's NetBurst architecture, and 400MHz FSB (100MHz quad-pumped) capabilities is expected for summer 2002 with cruising speeds in excess of 1.5GHz. Based on the Willamette core, the new

Celeron is expected to have 128KB of L2 cache.

The second half of 2003 will be the start of more exciting things to come with the 0.09-micron-based Prescott featuring more transistors (duh!), HyperThreading (first time on a desktop), and more cache. On the mobile side, Intel currently has the P4-M chip based upon the 0.13-micron Northwood core. The previously top-end Tualatin core used for the PIII-M will eventually be rebadged as a lower-powered, lower-cost, smaller-form-factor mobile Celeron for the rest of the year. However, a new mobile Celeron based upon the Northwood will likely enter the market before the end of 2002. Waiting until next year will open the door to Baniyas!

AMD keeps chipping away at the market share traditionally dominated by Intel. The company's Athlon XP is still the most compelling CPU on offer in terms of price and performance. Even though the current 0.18-micron Palomino-based Athlon XP 2100+ (clocked at 1.73GHz) can't win in terms of raw MHz, the superior IPC efficiency means that only a P4 2.4GHz pulls ahead in some applications, and AMD CPUs still offer the best overall value. The introduction of the 0.13-micron "Thoroughbred" core is due shortly. The die-size shrink from 0.18-micron should enable higher clock speeds and lower power consumption.

Before the year is out, AMD plans to introduce its first SOI-based 0.13-micron Barton core, which doubles the L2 cache from 256KB to 512KB. AMD will also be

Talking Future With Pat Gelsinger, Intel CTO

We caught up with Intel CTO Pat Gelsinger via his mobile phone one Tuesday evening as he was catching a plane to Japan. The following is an excerpt of that conversation. For the full interview, see www.smartcomputing.com/cpumag/jun02/gelsinger.

CPU: A lot of this development in different areas (BBUL packaging, CMOS Radio, TeraHertz Transistors, etc.) is happening in parallel. Is there some sort of hierarchical roadmap in the research department for when you think incremental changes will occur?

Gelsinger: We approach it in two ways. We have carved our research areas up into four major thrusts. One is the core silicon work, second is the core architecture platform work, next is the software layers, and then finally is the networking layer. Within each one of those four areas they have roadmaps of work that extend



Pat Gelsinger at the Spring 2002 Intel Developer Forum.

Bottom: Courtesy Intel Corporation.

Moore's Law: Beyond The Marketing

It gets airplay and magazine covers, but Moore's Law has always been more than a marketing slogan. Here is the word straight from a 20-year veteran of Intel R & D:

"If you go back and look at a particular aspect of the platform, for example, with disc drives, where we've needed to 'scale' the interface to take full advantage of the drive improvements, we've seen this 'scaling' occurring, more or less, in a periodic fashion. Regardless of the number, what we've really been able to observe is that there is a doubling over a certain period of time that for some reason or another has remained relatively constant. So where Moore's Law is typically applied to transistors of scale and frequencies of scale, we've seen this tendency in other aspects of the platform, as well. That is the motivation that drives us to look out in time and predict when we may need a new signaling technology, higher data rates etc., and to have the necessary platform ingredients in place. You may not hear the term 'Moore's Law' every day but it really does permeate just about everything we do."

Steve Pawlowski

Intel Fellow Director of Communications & Public-Community Technology

introducing its eighth-generation 64-bit Hammer architecture, manufactured on SOI technology. (We covered this on page 18 of the May 2002 *CPU*.) And the desktop market may get a dose of SOI 0.13-micron-based ClawHammer this year, as well. On the value side of the fence, the Duron will continue on at higher frequencies, thanks to a new 0.13-micron Appaloosa core, similar to the current Morgan core, but with a 266MHz FSB in the first half of 2002.

Achieving The Dream

So who is actually responsible for reaching the numbers that appear in such predictions as Gelsinger's vision of

2010? Intel's Microprocessor Research Labs (MRL; intel.com/research/mrl) develops key technologies for future microprocessors and platforms ahead of the product development groups. MRL is the leading organization setting the standard for microprocessor development three to 10 years in the future. Dr. Wilfred Pinfold is director of Marketing for MRL; here is how he characterizes his reaction to Gelsinger's prediction of 30GHz, 10 billion transistors, and 1 tera-instruction per second by 2010:

"Well, Pat is an interesting guy because he is both a very good businessman and he's also a very good technical guy, so he comes to those figures from a

good understanding of the technology, and I would say that reaching those figures will be a challenge and he knows they will be a challenge."

"... I look at numbers like this and I say 'Oh, we've really got our work cut out for us', but I do believe it's possible, and I believe the industry will gain an enormous benefit by us managing to stay with those numbers. There's a lot actually underneath those numbers that's not immediately visible. The technology that is required to hit those kinds of performance numbers can be applied not only to the performance vectors but also to the power vectors and to bringing a much wider range of products to [market].

"Not only will the technology that we need to develop over the next eight years (that he is talking about there for 2010) allow us to drive performance on the platform to meet the numbers that he's stating, but it will also allow us to give you laptops that run for multiples of days and systems that can run very cool ... and microprocessors that will be in small devices that you can have in your pocket [that give] immediate access to the digital world. You can use those technologies in a wide range of ways. So, Pat's given us a real challenge and that's what I'd expect from Pat. We will go and we will meet his challenge." **CPU**

by Alex Ross and Joan Wood

from two weeks to 10 years. The other thing we are trying to do, and this is where my unique job is, is to try and take a top-down view of that and find a relationship between different areas of research and pull those together into some of these longer-term visions that we are driving toward. Using the CMOS Radio and "Radio Free Intel" example, that combines four different areas of research at Intel. One was the MEMs work, one was the architecture of the CMOS work pulling off analog, another

was protocol work for roaming technologies, and another was software work in the area of software-defined radio. Part of my role is to pull those different areas together and then to try to assimilate out of that some of the longer-term views. So both the roadmap for each of those four areas and then a top-down view where we try to tie them together into some more compelling, longer term initiatives.

CPU: As Intel's new CTO, is there a certain amount of influence from your background in

Desktop and CPU development that exerts itself?

Gelsinger: Clearly my strong suit is in processors. That's where I spent the majority of my time, really. In the computer side of the house is where my natural technology and experience lies. So an awful lot of what I'm doing these days is getting smart on the communications side of technology. So over Christmas I had one of our optics experts come and spend literally a day giving me "Optics for Dummies." [laughs] So as CTO, my job is to crossover the

areas of wireless, of optics, of communications, and of computing, as well as the end user view and the network view of those technologies. So while I have good depth personally in the computing area, to do my job and do it well, it's essential that I get smart and build relationships with the key technologists both internally and externally and get those technologists to really complement the areas that I have as my own natural strengths. ▲

by Joan Wood

Video Cards

GPUs Of The Future

Consumer 3-D graphics hardware innovation is interactively linked with Microsoft's development of DirectX. New hardware features are rarely developed in complete isolation. Most concepts in 3-D have been floating around for years in research papers and discussion groups and are already available in the types of high-end software packages used to create animated films and special effects. The challenge for consumer graphics chip companies is getting these features onto silicon. Once this has been achieved, the companies then compete to have their particular implementation licensed by Microsoft for the next version, currently DirectX 9.

Future Features

A peek under the hood of publicly released DirectX 9 presentations gives clues as to what new features are expected in next-generation GPUs. Microsoft notes updated vertex and pixel shaders, increased support for higher order surfaces with two types of displacement mapping, added compression techniques for volume textures, and gamma correction. The notes also make recommendations of what features future hardware should support. They have indicated that future DirectX versions will include further expansions to gamma correction support and the integration of pixel and

vertex shaders, along with fog functions, into a general "graphics shaders" category.

Vertex and pixel shaders. When first introduced, a new feature is not often widely adopted by developers unless it simplifies some existing necessary operation or offers a dramatic improvement in speed, memory usage, or visual quality. As the feature matures, and implementation is improved over successive versions of the hardware, drivers, and DirectX, more developers integrate it into their graphics production pipeline. Most developers have not yet used pixel and vertex shaders extensively, so improving their usability could have a significant impact. NVIDIA and ATI are both very focused on advancing usability of their vertex and pixel shaders. It is quite likely they will succeed at this within the next generation of hardware.

D-Maps

Because "unannounced" Matrox hardware was used for Microsoft's demonstration of advanced displacement maps at the Game Developer Conference in March, it is a pretty safe bet that the Matrox method won at least part of the displacement maps slot in DirectX 9. ATI and NVIDIA are right in there, as well, with each company having its own preferences. Matrox talks about precomputed

D-mapping for characters and objects and sampled D-mapping with adaptive tessellation for terrain.

An oversimplified explanation is that while bump mapping makes an object appear to have bumps by painting lighting and shadows on the smooth object surface, displacement mapping actually changes the geometry of the object by adding the bumps. Both effects can be achieved using the same input data, so it is possible to imagine using them for LOD (level of detail) swapping so that an object would be drawn with "real" bumps only when you got close enough to notice that they were fake. As you retreat into the distance, bump mapping, which cuts down on geometry, would be substituted for the displacement mapping, and no one would be the wiser.

Displacement mapping techniques have been used for years in software but (apparently) now they can be done in hardware. Instead of storing an enormous 3-D model, a game engine can hand over a grayscale texture map for terrain height data (highest points represented by white pixels, lowest by black, and a range of gray for everything in between), a normals map, and specifications for a base mesh. The hardware will not only generate the mountain on the fly, it will generate the specified LOD appropriate for the scene, with lots of polygons for the mountain next door and simplified geometry when the mountain is far away. A very interesting possibility mentioned in the DirectX presentation

1/2MB Cirrus Logic. PCs often came with horrid graphics "decelerators" and no 3-D at all. The Cirrus Logic claimed to do 3-D and came with NASCAR racing but was just a great 2-D card, really.

VideoLogic PowerVR Series 1. This is the first consumer 3-D card. Ultim@te Race was cool; interesting tile-based rendering needed a fast CPU.

3Dfx Voodoo1 Orchid 3D 4MB. 3-D graphics hit the big time! Tomb Raider & Quake get hardware accelerated, and Glide was better than Direct3D.

Quantum 3D Obsidian 100SB-4440. The Obsidian has a Voodoo1 4/4/4 12MB layout and is the first SLI (scan-line interleave) product. It cost \$1,295 but could do 1,024 x 768! The Obsidian had possibly the longest PCB (printed circuit board) ever, with a spoiler on the end.

3Dfx Voodoo2 Diamond Monster 3D 12MB SLI. This was as fast as you could get, and it lasted 18 months! The Diamond Monster had three 3-D cards in all.

Quantum 3D Obsidian X-24. Why have two PCI slots filled when you can have 3Dfx Voodoo2 SLI on one board?

1993

1996

1997

1998

is a future D-Map tool that would take a high poly count model and convert it into a low poly model along with the necessary displacement maps. This could positively affect the adoption curve for developers.

NVIDIA's discussion of support for higher-order surfaces such as Bezier and Nurbs patches seems caged in the problems that the technology brings, so they may not be a near-term offering. However, NVIDIA does express keen interest in a Nurbs alternative: displaced subdivision surfaces. A combining of two long-known techniques, subdivisions surfaces and displacement maps, the main advantage over traditional methods appears to be a more accurate representation of the original model while maintaining very high data compression ratios. Alternately, ATI may be leaning toward a method that employs curved point normal triangles.

Future Hardware: Our Best Guess

Sometimes there is a fine line between what we know and what we can say we

know. If this were being written a week from now, we would have to lie and say we have no idea what XYZ graphics hardware company has in store because by then, we will have seen their hardware and signed an NDA specifically prohibiting us from disclosing that information until the release date. But right now we can truthfully say that we have no idea whatsoever and that all of this, except as noted, is pure speculation with nothing more than a decade of combined graphics and game developing experience and a week of trawling the Web's seedier neighborhoods behind it.

NVIDIA. NVIDIA's pattern is to release a new chip that does everything the last one did only better and with solid basic functionality for newly introduced features. The company designs each generation of chips to run with the previous generation drivers, so technically speaking, it has a shippable product straight from the fab (a manufacturing plant where chips are made). NVIDIA can further tweak drivers while OEMs buzz away building boards, and products can launch as soon as OEMs are ready. NVIDIA then spends the next six months on driver optimization, revealing more advanced functionality for new features as it enables them in the drivers. So expect an NV25 refresh with more polished drivers, improved performance, and a bump up in clock speed. Cost savings derived from using TSMC's 0.13-micron production process will also be key, as it should be live by then.

NVIDIA's acquisition of 3dfx and Gigapixel, which has thus far only been

revealed in small items such as LMA II (Lightspeed Memory Architecture II), may be in evidence at the release of DX9 later this year and could culminate in the integration of other acquired technology. Rendering will be bumped from 32-bit to 64-bit, thanks to more efficient use of bandwidth. The new texturing unit should be able to dynamically allocate resources to different pixel pipelines as needed. Vertex shaders will be pumped up with more registers, allowing for integration of more floating-point functions. Greater color precision, hardware shadow maps, higher-quality texture sampling, even better-looking AA is likely with NV30. Long term, look for NVIDIA to collect more expertise with the multimedia side of things through continued development of its integrated chipset offerings. Eventually everything will end up on one piece of silicon, except what the company cedes to Intel to keep on its good side.

ATI. ATI is currently playing the DX 8.1 game in second spot with its R200-based RADEON 8500, but looking ahead, the R300 is shaping up for DX 9. In line with the DX 9 spec, the R300 promises game developers easier utilization of fully programmable Pixel & Vertex shader functions, which will hopefully trickle down into much more content. ATI will likely improve performance under Antialiasing with a better AA engine, which currently severely lags behind NVIDIA's GF4 Ti. Rumors point toward a stochastic (random pattern) type of AA that will utilize a 5 x 5 matrix.

Worthy Of Note

- Rendition's VQuake
- 3dfx Voodoo 2 dominates for 18 months
- Matrox 2-D innovations such as DualHead
- ATI brings 3-D gaming to the mobile PC
- 3dfx tries to push FSAA with VSA-100
- S3 for introducing S3TC (texture compression)
- Bitboys Oy for having the best press releases
- NVIDIA picking up 3dfx & Gigapixel on the cheap

3dfx Voodoo3 3500. This card was fast, but 3dfx still didn't have 32-bit.

NVIDIA TNT2 Ultra. NVIDIA overtakes 3dfx. The TNT2 had great 2-D/3-D all in one and was the end of the line for 3-D only. Later, GeForce 256 introduces T&L to the industry. It started with SDR RAM and moved to DDR. NVIDIA gets aggressive and innovates for DX7.

NVIDIA GeForce2 GTS. The six-month product cycle churns with more pipelines, more features, and more speed for GF2 GTS.

NVIDIA GeForce2 Ultra. At the end of the year, a refresh means a faster clock speed and now with 64MB!

NVIDIA GeForce3. This card was expensive but had revolutionary Pixel & Vertex Shaders. NV20 gets a refresh with GeForce3 Ti 500.

ATI RADEON 8500 128MB All-in-Wonder. Gamers toy with the feature-rich GF3 challenger, the ATI RADEON 8500, by year's end.

NVIDIA GeForce4 Ti 4600. The GeForce4 has fantastic FSAA performance and 128MB. The NV30 from NVIDIA and the R300 from ATI are on every gamer's Christmas wish list.



1999

2000

2001

2002

The R300 is being developed by ATI's acquired ArtX team, which conjured up the "Flipper" 3-D chip for the Nintendo Gamecube. Flipper, abnormal for bottlenosed dolphins, sports 3MB of embedded SRAM. The embedded DRAM is beneficial, as it provides even more bandwidth, has been used on mobile RADEONs already, and could become a decent way to tackle bandwidth

constraints. MPEG-4 acceleration also seems likely.

Matrox. Since introducing "EMBM" (environment mapped bump mapping), Matrox has taken a relative backseat in the 3-D gaming market, letting the rest fight it out and instead focusing on its bread and butter: 2-D. The corporate market aside, most gamers haven't used a Matrox card in a very long time. 3-D performance has been slow, especially under OpenGL, and the feature set is rather behind the rest of the pack. However, the advent of DX9 is promising to change that with rumors already rampant regarding a next-generation chip dubbed "Parhelia," also readied for OpenGL 2.0 support. The company is banking on hardware acceleration for displacement maps. Matrox has been evangelizing displacement maps at insider trade shows such as Microsoft's Meltdown for the best part of a year.

In order to compete in the gaming market and be fully DX9 compliant, Matrox will obviously have to introduce a programmable T&L (pixel & vertex shader) unit on-chip for the first time. That much is a given. It looks likely that a product with 128MB of high-speed DDR-SDRAM will provide roughly 20GBps of bandwidth. In line with Matrox's 2-D strengths, DualHead will be standard. The architecture will be completely new and not simply a revamped G550, which was actually based on the old G400. Matrox doesn't have the same muscle as NVIDIA or ATI when it comes to manufacturing,

but in order to hit a "rumored" 300+MHz core speed, a chip based upon a 0.15-micron process is likely.

Other Hopefuls?

Remember when Creative Labs was king of the OEMs when it came to 3-D graphics? It left the business a few years back, but with the acquisition of old graphics stalwart 3DLabs, the company is obviously readying a comeback. At a presentation, Creative CEO, Simon Wong, boasted that the architecture being used will be the first tera-op processor with 76 million transistors at 170 gigaflops and 64-bit color support. The company's product is rumored to be shipping around Christmas.

Everything eventually migrates to silicon. Features migrate downward from non-real-time animation software on workstations and high-end real-time simulation hardware down to the consumer PC. Once there, it migrates further. From very expensive first-gen cards to second-gen midrange offerings to low cost, low power/low heat mobile and finally down to integrated solutions; features, performance, and price finally converge on the motherboard. On this account, don't discount Intel for the long-range plan.

And don't rule out OpenGL. The "design by committee" progression of the Silicon Graphics API has a very different development model to DirectX. It generally lags behind new hardware implementation developments by months or even years, but that is mitigated somewhat by OpenGL's flexibility to add GL extensions: custom bits of code that each programmer creates to implement their own features. Another potential advantage to the OpenGL model is that many voting members of the OpenGL Architecture Review Board are the same hardware companies battling it out over market share and DirectX credits. But they eventually have to sit down and play nice with each other to get features implemented. Instead of rushing to be first out the gate, they can exchange ideas, discuss, and investigate to come up with the best overall solution. It will be interesting to look back in 10 years to see which method has had more enduring results. **GPU**

by Joan Wood and Alex Ross

Buying Tips

New features are introduced in products that become outmoded before games that use those features are released. So by the time you want to play Doom III with everything turned on, NVIDIA will be on its second or third rev of products that support those features. This is great for consumers because bugs have been worked out, and the price has come down. Let the guy next door be the first to pay through the nose for the privilege of watching next-gen demos. Use the cash you save to buy more games that play perfectly well on the second-generation card you just got for an amazing price.

1 Wait until there is a game that you really want to play with every feature turned on before purchasing the hardware that supports those features.

2 Look for the least expensive hardware that supports the features you actually want. Often that will be at least a generation older than the current "state of the art" board.

3 Even then, look for a card that also offers improvements that affect all your games, such as antialiasing, extra memory, or high throughput.

4 Research it yourself. Use multiple sources for reviews and get an idea of what you want before going to the manufacturer's Web site and checking availability and its listed prices. Then go to Price Watch or one of the other price search engines to get the online prices from specific dealers.

5 When you go to a retail store, don't depend on the salesperson knowing much about specific products.

Bloopers

- NVIDIA almost going bust with NV1
- 3dfx trying 2-D/3-D too early in 1997 with awful Voodoo Rush
- 3dfx calling Banshee2 a Voodoo3
- 3dfx flushing millions in profits on world's worst marketing campaign
- NVIDIA's Voodoo Bicycle
- Metabyte claiming TNT2-SLI technology
- ATI taking forever to actually release the Rage Fury 128
- S3 claiming T&L on the box



MORE VIDEO TIPS FROM SCOTT & ERIC™

“Music is a tried and true way to add drama and impact to your videos. As you can see here.”

Before music.



After music.



DUN-dun-DUN-dun-DUN-dun... What would the shark footage in *Jaws* be without that, huh? Or what if they'd used “It's A Small World” for the shower scene in *Psycho*?

Clearly, music can make a world of difference—especially in your home videos. And Pinnacle Studio™ version 7 makes it amazingly easy to jazz up your videos with music—your own or choose from a built-in library of songs and sounds—along with hundreds of special effects,

titles, and transitions. Then burn 'em to CD or DVD with Pinnacle Express™. Now, along with preserving memories, you're creating tightly edited cinematic works of art. It's fast. It's a blast. And it's made possible by Pinnacle Systems, the folks who've won 8 Emmys for Technical Achievement. So, whether it's “Home On The Range” or an old Def Leppard tune, adding music is just one of the ways Pinnacle can bring your videos to life.

Get a \$25 manufacturer's rebate.*



Buy Pinnacle Studio version 7 and Pinnacle Express DVD/CD authoring software and we'll send you a check for \$25. (Sorry, the free Porsche was a no go. — Scott & Eric) For details and more tips, visit www.pinnaclesys.com/memory.

P PINNACLE®

Available at:



COMPUSA



MICRO CENTER

*Buy ANY Pinnacle Studio version 7 product AND any Pinnacle Express product between 2/15/02 and 4/30/02 and fill out our Pinnacle Express/Studio coupon available at www.pinnaclesys.com/memory. Return completed coupon, dated sales receipt, and the official UPC codes from the boxes of both products to the address on the coupon to receive your \$25.00 rebate check. Rebate request must be postmarked no later than 5/31/02. Allow 6-8 weeks for delivery. ©2002 Pinnacle Systems, Inc. Pinnacle Systems is a registered trademark and Pinnacle Studio, Pinnacle Express and 'More video tips from Scott & Eric' are trademarks of Pinnacle Systems, Inc. All other trademarks are the property of their respective owners.

Buses & Chipsets

Coming Soon To A PC Near You

Most of the chipset introductions for the remainder of 2002 will represent less of a next generation than they will a maturing of the current crop. AMD, Intel, VIA, SiS, NVIDIA, and now ATI will be aiming most of their new wares at addressing ever faster memory chips (especially the DDR333 and 400 flavors) and integrating more functionality onto the motherboard. Later in 2002 and early in 2003, when AMD releases its x86-64 Hammer architecture and Intel moves to even faster Pentium 4s, we will see the bus wars heat up and the architectures start changing noticeably.

Pentium 4s For All

Intent on keeping P4 buyers riding its own bus, Intel is issuing as many as six new chipsets this year. Aside from several integrated products for the bargain market, power users should look for the Rambus-based (RDRAM) 850E upgrade to the current 850 standard. This will add the 533MHz FSB that upcoming Pentium 4s will exploit and support for newer PC1066 RDRAM modules. Intel continues to deepen its newfound support of DDR memory, as well. New chipsets for the Pentium 4 system will support a 533MHz FSB, PC2100 DDR and USB 2.0 (dubbed the "Brookdale-E").

In mid-2003, we should see the Springdale chipset, which will run Pentium

4 Northwood and Prescott processors and is reported to include support for the next-gen spec for DDR memory DDRII SDRAM, AGP 8X and a 533MHz (133MHz quad-pumped) FSB.

SiS will be trying to supercharge the P4 in a number of ways. The company will follow up its recent debut of the SiS645DX, which supports a 533MHz FSB and DDR333/400, with the SiS648, which will add support for AGP 8X. A later version, the SiSR648, will support 533MHz RDRAM (PC1066), making SiS the only non-Intel manufacturer to design chipsets around Rambus memory. The company is also introducing a Southbridge chip (the SiS962) that will integrate IEEE 1394A and USB 2.0 support.

Server scene. The long-promised 870 server chipset from Intel may see the light of day by the end of 2002. The 870 is designed to support the new Xeon and the McKinley (the successor to the Itanium) with DDR memory. The latest word is that it will support 512 processors at once. The 870 chipset will also be used for the Madison and Deerfield processors, which should be out sometime in 2003.

Support For Athlon, Hammer & Beyond

The next major leap for AMD's homegrown chipsets will come late in 2002 with the Hammer family of processors, which

will ride on a wholly new AMD 8000 family of chipsets. We've already seen the first few iterations of this series. (See pages 18 and 19 of May 2002 *CPU*.) For starters, the architecture carries a new I/O chip, AGP 8X support, and PCI-X. Although details have not emerged, all the usual AMD motherboard suspects, VIA, Ali, and SiS, are preparing similar chipsets to support the Hammer line when it arrives.

Following up on its new KT333 chipset for Athlons, VIA will support DDR400 memory chips in the KT400 later this year. Its Hammer family support will also come in the fourth quarter with the K8HTB, including AGP8X and a 64-bit PCI bus.

The second-gen nForce chipsets from NVIDIA incorporate the line's existing feature set (integrated GeForce2 graphics, 5.1 channel audio sound, and faster memory and HyperTransport I/O) but will add DDR33 memory in the 6220-D and 615-D models. Look out next year, however, when the "Crush 17 and 18" chipsets up the video ante with a built-in GeForce4 MX and probably DDR400 SDRAM support.

ATI is jumping into the market with sets that support Intel and AMD processors across several platforms. The RADEON IGP 330/340 and 340M will power P4s and mobile Pentiums, while the RADEON IGP 320 and 320M are for AMD and Durons on desktops and notebooks. (See page 20 of May 2002 *CPU*.) All models will surround a RADEON 7000 core, but by early 2003, ATI may have new chipsets that use the upcoming R200 graphics processor.

Notebooks

VIA Technologies is bringing DDR266 memory speeds to the notebook with its

The Original IBM chipset for the PC supports Intel 8088 processor running at 4.77MHz, cassette port, floppies, 64KB of on-board memory, and five 8-bit expansion slots.

Intel's 4004 microprocessor has a 4-bit bus.

The ISA standard expands to include a 16-bit bus and a 16.6Mbps data transfer speed.

Intel introduces the 430FX chipset to support its new Pentium processor, including EDO RAM and PCI 2.0 support.

1971

1981

1984

1992

ProSavageDDR PN266T for Pentium III-M systems, which is also 30% lower voltage for longer battery life. Meanwhile, as mentioned above, ATI is making a play for the notebook market, as well, with chipsets for mobile Pentiums and AMD processors that integrate the RADEON graphics core, naturally. The heralded NVIDIA nForce chipset will also move to the notebook market in the last quarter of 2002 with the "Mobile Crush" chipset.


Although not confirmed as part of any announced chipset, the industry expects 802.11b wireless and USB 2.0 support to become standard in notebook chipsets in 2003. Toshiba, Dell, and IBM have already introduced models with 802.11b cards integrated within the notebook itself, and Gartner Group projects that 90% of mobile PCs sold will have WiFi built in by 2005.

AGP: The Serial

The long-promised AGP 8X slot is scheduled to be in a number of the chipsets we've already mentioned here by the end of 2002. Now dubbed AGP 3.0, the spec calls for a 32-bit wide parallel bus running at 533MHz. The effective bandwidth will be 2.1GBps, twice the throughput of the current AGP 4X standard.

But 8X may be the last X AGP sees. A consortium of chipset and PC makers,

A NEW AGP



By the end of 2002, many new chipsets should feature the new AGP 8X. With an effective bandwidth of 2.1GBps, AGP 8X represents a significant speed increase from AGP 4X. After that, look for serial AGP, which could provide a total bandwidth of about 10GBps but shouldn't be available until around 2004.

including Dell, Intel, IBM, and NVIDIA, is working out the next graphics implementation for motherboards, which will take the video card off of the AGP bus and put it onto the 3GIO serial bus. No longer running in parallel mode, this "serial AGP" spec calls for a bus that can scale from 1 to 32 bits and may be able to reach bandwidths of 10GBps. Most likely, the spec will not be backward compatible with current AGP video cards. We won't be seeing "serial AGP" until sometime around 2004.

Take The Next Bus

Intel and AMD will move their titanic battle over CPUs to the basic system bus

within the next two years or so as each rolls out competing chipsets that claim to change the ways in which individual motherboard chips communicate with one another. Intel is banking on 3GIO, which is designed to accelerate communications among a motherboard's subsystems and its add-in cards. Intel is planning an "Evo-Revo" architecture design that will let users plug everything from MP3 players to hard drives in and out of their systems as easily as you put game cartridges into a Game Boy Advance. Intel is hoping to see 3GIO launch in late 2003 and early 2004. (Learn more about 3GIO in this month's White Paper on page 48.)

AMD is rallying support for its alternative HyperTransport architecture, in development since 1997. Its upcoming Hammer CPUs will support this new approach, which promises to decrease the number of chip-to-chip buses overall on a motherboard and allow chips from different vendors to work with one another more smoothly. The specs have raised eyebrows with the promise to eliminate system bottlenecks by accelerating chip-to-chip communications 40 times and up to 12.8GBps. AMD's HyperTransport Consortium of allied vendors has grown to more than 70 members, in part because manufacturers see this design as a way to mix and match subsystem chipsets more effectively and so differentiate their product lines.

Which is what all of us really want: even more chipset variations to confuse our next PC or notebook purchase. **COU**

by Steve Smith

The industry adopts the PCI standard, a 32-bit bus that runs at 33MHz and has a 133MBps bandwidth. PCI remains the industry standard today.

The Intel 440LX chipset is the company's first to support the AGP graphics bus.

Intel's 440BX Pentium II chipset dominates the market, bringing 100MHz bus clock speeds.

PCI-X finally begins appearing in hardware. A 64-bit PCI-X bus can carry 1.1GBps.

AMD introduces its HyperTransport technology. Intel announces its 3GIO standard. The PCI-SIG announces support for 3GIO.

The Arapahoe Work Group announces the first draft of the 3GIO specification. Chipsets from Intel, NVIDIA, ATI, and VIA are expected to support AGP 8X.

Hardware featuring 3GIO is expected to appear in late 2003 or 2004.

1993 1997 1998 2001 2002

The Bleeding Edge Of Software

by Warren Ernst

Inside The World Of Betas

Fortune-telling becomes an exact science when it comes to working with betas. This month we bring you six glimpses into the future of software.

iSiloX 3.11 beta2

For owners of Palm OS-powered PDAs, life would be beautiful if you could easily snatch Web pages and stuff them into your device. Easily converting portable graphics and office files would also be nice. Well, direct your attention to iSiloX, a combination offline browser, file converter, and "site sucker." In short, iSiloX makes highly compressed files for your Palm OS device that you view with the iSilo reader.

At its heart, iSiloX converts documents from plain text or HTML into highly compressed files that can contain graphics, text, tables, and links—almost like a small, self-contained Web site. The application can also follow links and retrieve text and graphics from Web sites, building a readable file as it goes. In addition, it

can grab data located off a targeted site if the first site links to it.

iSiloX also does a terrific job compressing data. For example, I compressed a cocktail recipe Web site containing about 1,200 drinks into one 240KB file, complete with thousands of hypertext links. Depending on the PDA, graphics may convert to different color depths.

Version 3.11 beta2 is largely a bug-fix version, and it seemed largely free from crashes and hang-ups. Conversion speed is still an issue, as a large site can take several minutes. The sheer number of options helps get most of a site into the PDA, and a well-written Help file makes it almost simple. If you use a Palm OS-powered device, check out iSiloX, pronto.

Official product name: iSiloX
Version # previewed: 3.11 beta2
Publisher: DC & Co.
Developer and URL: DC & Co;
www.isilox.com
ETA: Q2 2002
Why you should care: Portable Web documents made easy and compact.

OpenOffice Release 641d

News about StarOffice might seem, well, familiar. About four months ago, we highlighted the StarOffice 6 beta. Initially, StarOffice was to be a free, drop-in replacement for Microsoft Office 2000 available in Q2 2002. Now, Sun Microsystems apparently will add customer support and extra features and charge about \$100 or so for the package to meet the demands of users who want a Microsoft alternative that also offers customer support.

A free open source version is still available, however, at OpenOffice (www.openoffice.org). Sun is a bit cagey explaining the differences between StarOffice and OpenOffice. After using both for some time, however, for the end

user OpenOffice is like StarOffice, but without a database module or bundled templates or fonts, fewer import filters, and no official support from Sun. (For developers, OpenOffice is very different—much like Mozilla is to Netscape.)

By and large, this is a good thing. OpenOffice contains word processor, spreadsheet, presentation, and drawing apps, each with about 85% of the functionality of Word, Excel, and PowerPoint. OpenOffice opens and saves these Microsoft files automatically and maintains formatting

and formulas throughout. The Windows version of OpenOffice looks and feels like a Windows app, rather than a Linux or Java program ported over.

Version 641 is the final beta version and is intended to catch the last bugs rather than introduce new features. However, the version is stable enough now to use all day every day.

Official product name: OpenOffice
Version # previewed: 641d
Publisher: The OpenOffice Organization
Developer and URL: The OpenOffice Organization; www.openoffice.org
ETA: Q3 2002
Why you should care: A free, high-quality competitor to Microsoft Office.

XPlay Technology Preview 5

There are things you just shouldn't mix, such as grape jelly and taco sauce. Such was the case with the iPod and a Windows-based PC. We might love each separately, but they seemingly weren't meant to intermesh—until now.

XPlay for Windows provides full iPod support in Windows—almost. The app is very much a beta (or even an alpha) product, and the readme is fraught with dire warnings about reformatting steps should something go horribly wrong, including fundamental FireWire problems with various firmware and Windows flavors. That said, XPlay worked fine for me, doing nothing terrible to a friend's cherished iPod (whew!).

Using XPlay is straightforward, and not. Using Mediafour's own MacDrive technology, XPlay makes a connected iPod appear as a 5GB drive in Explorer and My Computer. Within this is an XPlay Music

icon, and within *that* is your music, sorted by Artist, Song Title, Album, and Playlist, like Apple's iTunes. Double-click an icon and the file starts playing in Media Player. Dragging and dropping music files doesn't work, which is somewhat by design so you can't use the iPod to illegally transfer music to another computer (presumably at Apple's request). You can drag and drop nonmusic files into the iPod.

XPlay also adds functionality to Media Player 7.1. Media Player recognizes the iPod automatically once XPlay is installed, letting you create and play playlists, view visualizations, and copy ripped files. Xplay isn't for everyone, but despite its small problems, it is amazing.



Official product name: XPlay
Version # previewed: Technology Preview 5

Publisher: Mediafour
Developer and URL: Mediafour;
www.mediafour.com

ETA: Q3 2002

Why you should care: Windows and the iPod: two great tastes together—finally.

NetCaptor 7.0 beta 4A



Official product name: NetCaptor
Version # previewed: 7.0 beta 4A

Publisher: Stilessoft

Developer and URL: Stilessoft;
www.netcaptor.com

ETA: Q3 2002

Why you should care: Nearly all of Opera's features and IE's rendering engine makes for a superior browser.

What do you get when you combine the guts of Internet Explorer with touches of the brilliant interface and feature set of the innovative and compatible Opera? You get NetCaptor. If you really like Opera but hate how some Web pages don't look right, NetCaptor may be your dream browser.

NetCaptor makes use of IE's browser engine. As a result, it renders Web pages exactly like IE, which most developers write for, so you get virtually no on-screen glitches. Even complicated plug-ins, such as Windows Media Player and RealOne, work without modification.

However, surrounding the Web page is a suite of features any Opera user wouldn't want to live without. The most important are Tabs that let a single window hold many Web pages. There's also a built-in pop-up window killer, interactive JavaScript, ActiveX, cookie management, automatic translation

and what's related menus, Alexa and Google reviews, cached screen viewing, and even a hookup to Anonymizer.com.

Perhaps the most unique feature is Groups, collections of URLs that open in their own Tabs simultaneously with a click. For example, the News Group automatically loads CNN, Excite News, The New York Times, Los Angeles Times, PR Newswire, and Newspapers.com. There are prebuilt Groups for shopping, employment, searching, and software downloads. Making new groups is extremely simple.

NetCaptor lacks Opera's ability to wave the mouse pointer around for back, forward, and other commands, nor is there Opera's amazing array of keyboard shortcuts. The beta also lacks Opera's legendary stability. These negatives aside, NetCaptor is worth a serious look.

Infinite Loop

Ancient Art & Computers Collide

Forget wing-flapping cranes. Coupled with CAD, origami is heading into space and elsewhere. In the early 1990s, engineer Robert Lang designed TreeMaker, a program that uses an algorithm to convert basic sketches into origami patterns. Recently, Lang helped the Lawrence Livermore National Laboratory design the

hinged-, folding-paneled lens of a 100-meter-in-diameter space telescope. Back on Earth, engineer Rainer Hoffman and origami expert Tomoko Fuse combined CAD and origami techniques to fold vehicle air bags for safer deployment. And a program called Molecular Origami helps scientists create protein-folding models.



Vamp 3.09 Beta

If you're looking for a new way to use your multimedia files or you want to combine them into animations and screen savers or share them with friends, Vamp may be what you're looking for.

Vamp is a graphics file viewer/animation displayer/MP3 player/file sharer/media downloader, and more. That sounds like a lot for a single program to handle, and it may be. The application plain freaked out on many AVI files I tried to display and on nearly every MP3 in my library.

What does work in the program is still interesting. When installed, Vamp automatically searches your drives for photos and pictures, videos, songs, and more and indexes them. It blasted through about 20,000 files on my test system in less than two minutes. It also connects to the Vamp Web site and offers to retrieve multimedia sets that other Vamp users have created and passed to their free Vamp accounts. Obviously, you can do the same

with your own free account and contribute to the astonishingly huge online library.

You can combine these files into online greeting cards, presentations, calendars, wallpaper, and screen savers, complete with animated transitions, specialized timings, and MP3-based audio tracks. Vamp's somewhat cluttered interface uses drag and drop to arrange material, much like video-editing software. The app is easy to use thanks to helpful labels and numbered steps, and it can also print and e-mail a resulting package directly to someone.

Vamp also offers a collection of online games, which are primarily simple Shockwave-based affairs. However, having a collection of games all in one place is a convenient distraction.



Official product name: Vamp
Version # previewed: 3.09 Beta
Publisher: Interactive Music
Developer and URL: Interactive Music; www.imvamp.com
ETA: Q3 2002
Why you should care: A multimedia viewer and file sharer in one.

RunIt! 3.0.0 beta2

There wouldn't seem to be a whole lot of room for improvement in the Windows Run dialog box, but the developers of RunIt! don't agree. RunIt! adds a whole passel of features to that dialog box, plus a few power-user features you probably never knew you wanted.

The application turbocharges the Run dialog box by first scanning your system for executable applications. Then, as you type a program or URL to run, it filters the list for you to choose from. If you have any arguments, there's a field for them—no more using quotation marks. It also keeps a History of the programs you've run, with the complete paths.

Although RunIt! doesn't automatically replace the standard Run dialog box, it does the next best thing by adding a set of icons to the Taskbar for instant accessibility. The program also keeps track of any application you run via any method, be it from the Start menu, a Desktop icon, or from another program. These are kept in a separate History list, a maximum 100 items, also available from the Taskbar. There's also a Terminate Application list, which

doesn't require scrolling and terminates programs without prompting (two big improvements over the system default).

The program has a high degree of configurability. You can instruct it to ignore certain apps, apps that launch from certain folders, or apps from certain drives. You can also set shortcuts in a Quickstart menu.

The only problems I found with the beta2 version (beta version 3 was released as we were going to press) related to the Help file not opening at every Help button that's offered and an almost unnoticeable slowdown in overall system performance on a slow system. But those are OK trade-offs for the extra functionality RunIt! adds.


Send Us Your Betas

Know of software in the beta stage that's deserving of some attention? Let us know. We'll take a look at it and possibly give it a go around. Send your prospects to bleedingedge@cpumag.com.



Official product name: RunIt!
Version # previewed: 3.0.0 beta2
Publisher: Keleos Software;
www.keleos.com
Developer and URL: Keleos Software;
www.keleos.com
ETA: Q2 2002
Why you should care: Turbocharge the Windows Run dialog box.

Virtual Competition: Connectix Virtual PC 4.3 vs. VMware Workstation 3.0



Virtual PC 4.3

\$199

Connectix

www.connectix.com



Having more than one OS installed on your computer is convenient but dealing with multiple drive partitions and temperamental boot loaders isn't. Virtual PC 4.3 and VMware Workstation 3.0 offer no-fuss alternatives. The apps create virtual machines that run multiple OSes—including DOS, Linux, and Windows versions—under your present Windows or Linux installation.

Connectix Virtual PC 4.3

Connectix is a long-time player in the emulator field, and the company's expertise shows. Virtual PC 4.3 is a polished, feature-filled app that works on a wider array of Windows systems than VMware Workstation, including Windows 98SE/Me/NT/2000/XP. The program also supports more guest OSes (the OSes that run virtual machines), including NetWare, Solaris, OS/2, DOS, Win3.1 and later, and Linux.

I tested Virtual PC with WinMe, Red Hat Linux, and FreeDOS. All three worked admirably. In some instances, Virtual PC is perkier than VMware, especially for

full-screen text interfaces. A handy control panel with miniature versions of each PC's screen lets you switch between guest OSes. Virtual PC's guest OSes are fast enough to run office applications, but like VMware, complex games are out of the question.

For Windows-based guest OSes, you can install Virtual PC Additions, which add such helpful features as folder sharing between guest and host OSes and the ability to drag and drop files to copy them between OSes. An arbitrary Desktop resizing feature lets you quickly change the resolution of a Windows guest PC screen simply by stretching or shrinking the Virtual PC window.

Connectix offers OS Pack Add-ons, which are virtual drives with preinstalled OSes. These are analogous to VMware's Guest OS Kits, but Connectix offers a wider variety of ready-to-run OSes, including Win98/Me/2000/XP Professional and Home. I tested the WinMe OS Pack. After some initial difficulty uncompressing its

enormous Zip file, I was running WinMe on a virtual machine without the tedium of installing it. OS Packs cost \$149 to \$299.

A few of Virtual PC's features edge out similar functions in VMware. For instance, both programs offer a full-screen mode, but Virtual PC lets you lock users in a full-screen mode, requiring a password to regain access to the host OS. Each guest PC includes a built-in Virtual Network Computing server, letting you operate the guest OS via remote control with a standard VNC client (www.uk.research.att.com/vnc).

Virtual PC does fall short in hardware support. It handles printers, serial ports, CD-ROMs, and floppy drives, but it can't handle USB devices that require special drivers. Our digital camera was invisible to guest OSes. (USB keyboards, mice, and floppy drives are supported, however.) Virtual PC also lacks some of VMware's advanced hardware functions, such as SCSI drive emulation and unlimited virtual drives (Virtual PC's OSes are limited to three.)

Emulate A Mac On Windows?

Vmware and Virtual PC work wonders for running Linux, DOS, Windows, and other PC-based OSes in a virtual machine, but what if you want to use Macintosh software on your PC?

That's more complicated. Because Macintosh uses a Motorola CPU and proprietary ROMs, a more complex emulator is necessary. The state of the art of Macintosh emulation lags far behind that of PC virtual machines. Many Mac emulators are more of a novelty than useful productivity tools.

However, options include SoftMac (www.emulators.com/softmac.htm), Basalisk II (www.uni-mainz.de/~bauec002/B2Main.html), Executor (www.ardi.com/executor.php), and The vMac Project (www.vmac.org). Don't expect these emulators to run OS X. Many can't even run PowerPC software (they're limited to ancient 68000-based code and obsolete versions of MacOS).

Some of these programs include a reverse-engineered version of the Mac's ROMs, the toolbox

that is the heart of the Mac interface. These reverse-engineered ROMs are free and legal but aren't as reliable as the real thing. If your software doesn't include these, you'll have to provide a software copy of Mac ROMs (which will certainly raise the hackles of Apple's lawyers) or install actual Mac ROMs on a PC Card.

The bottom line: If you occasionally need to run Macintosh software, get a used Mac from the classified ads or eBay. ▲

Where VMware is available for Windows and Linux, Virtual PC is available for Windows and Macintosh. Virtual PC's hard drive files are compatible with the Mac version, so you can move virtual PCs between platforms.

At \$199, Virtual PC is \$100 less than VMware. A 45-day trial version is available for download from www.connectix.com. The program has a bit more polish than VMware Workstation, and extra touches



VMware Workstation 3.0

\$299

VMware

www.vmware.com



make it a great choice, even if you're not a computer geek. The only major drawback is a lack of support for USB peripherals.

VMware Workstation 3.0

VMware Workstation 3.0 runs DOS, Linux, any Windows version, and FreeBSD. You can even run several OSes simultaneously. VMware says the program doesn't work with OS/2; Solaris and NetWare may work but aren't officially supported. The host PC you run VMware on must be running Linux or WinNT 4.0/2000/XP.

I tested VMware with WinXP Pro, Red Hat Linux, and FreeDOS. VMware worked commendably with all three; the OSes remained blissfully unaware they were living in captivity. There is some overhead in running one OS within another, but virtual computers running in VMware are more

than fast enough for word processors, Web browsers, and other business apps. All-text modes were noticeably sluggish, especially during scrolling.

VMware's hardware emulation is particularly strong, offering four networking modes for guest OSes, including NAT (Network Address Translation) and bridged networking. VMware uses virtual hard drives to store guest OS data. The virtual hard drives can emulate IDE or SCSI, and you can control the bus parameters of virtual drives (deciding if each is a master or slave drive, for example). Each guest OS can have any number of virtual hard drives and can access files on physical drives, too.

Support for many USB devices is also included. I was able to download images from a digital camera without fuss using WinMe as a guest OS. VMware doesn't support USB modems or devices that stream audio and video.

When you install Windows or a graphical Linux system within VMware, the guest OS display is set to an agonizingly low

I tested VMware with

WinXP Pro, Red Hat

Linux, and FreeDOS.

VMware worked

commendably with all

three; the OSes remained

blissfully unaware they

were living in captivity.

resolution and bit depth. Installing VMware Tools (a special set of graphics drivers) is the solution. This is an easy, one-click operation under Windows. Installing the Linux version wasn't as easy; uncompressing the

archive and running the install script is a challenge when you can see only a quarter of the screen.

You can install guest OSes from CD, floppies, or ISO images you download. Alternately, buying an optional Guest OS Kit gives you a ready-to-use virtual drive with a pre-installed OS. The only Guest OS Kit currently available is for Win2000 (\$199), but a WinXP version is planned.

VMware costs \$299. You can try it for 30 days free by downloading it from www.vmware.com. VMware Workstation doesn't match Virtual PC feature for feature, but its advanced hardware emulation and USB support make it an excellent choice. **CPU**

by Kevin Savetz

What They Can & Can't Do

	Workstation 3.0	Virtual PC 4.3
USB Device Support	Yes	No
Full-screen Mode	Yes	Yes
Ability To Lock Users In Guest OS	No	Yes
Ready-to-run OS Packs	Win2000 only; WinXP is planned	Win98/Me/2000/XP Pro and Home
Officially Supported Guest OSes	Win3.1 and later, DOS, Linux, FreeBSD	Win3.1 and later, DOS, Linux, NetWare, Solaris, and OS/2
OSes Supported	WinNT/2000/XP and Linux	Win98/Me/NT/2000/XP and Mac
Free Trial	Yes	Yes
Cost	\$299	\$199

by Kevin Savetz

Easy CD-DA Extractor 4.6.9

Easy CD-DA Extractor rips CDs, converts files between audio formats, and burns files to CD, but Extractor probably won't revolutionize your digital audio collection.

The application can encode music in a variety of popular and lesser-known formats, including MP3, Ogg Vorbis, and several flavors of WAV, AIFF, TwinVQ, and Monkey's Audio.


I ripped several CDs to create MP3s and other audio files, with the resulting sound quality being fine. For MP3s, Extractor lets you choose Faster or Higher quality encoding from 8Kbps to 320Kbps, including variable bit-rate encoding. But some files that Extractor created confused Windows Media Player into believing the total length was less than a minute, even though Media Player was able to play the entire 4-minute song.

Easy CD-DA Extractor is compatible with the freedb database and includes the ability to submit disc information to freedb. You can also play CDs directly from Extractor, but I had mixed

results. The program became so unresponsive on WinXP when playing a track I could barely make it switch tracks. Pressing the Stop button inexplicably locked up the program. I had no such problems using Extractor with Win2000.

The package also includes Easy Audio File Converter, an efficient tool for converting audio files between formats via a drag-and-drop interface. There's also Easy Audio CD Creator for burning audio CDs from MP3 and WAV files. Creator is decidedly lacking in frills; it will create audio CDs but can't burn data-only MP3 CDs.

If you get stuck, you'll be glad to have trial and error on your side because the online help consists of only a few terse Web pages of limited troubleshooting information and unhelpful descriptions of obvious functions. Tech support didn't respond to our queries. Easy CD-DA Extractor has potential, but it's unreliable and its features are limited compared with other audio tools. ▲



Easy CD-DA Extractor 4.6.9

Shareware; \$24.95

Jukka Poikolainen

www.poikosoftware.com



Fanix Software As-U-Type 1.2

Your word processor has a spell checker, naturally. But that doesn't do you much good when you're writing an e-mail message, posting to a message board, or composing a flame in your Usenet client. As-U-Type will improve your spelling no matter the applications you're using.

The program watches keystrokes as you type, immediately correcting any typing mistakes. Real-time spell checkers aren't new, but As-U-Type is highly customizable, yet practically invisible and doesn't interrupt your work.

If you make a mistake, such as typing "does't" instead of "doesn't," As-U-Type silently corrects the error. If it has no clue what you meant, it politely beeps, although its guesses are usually accurate. If it misses, As-U-Type learns from its mistake. For instance, if your fumbling fingers type "peaker," the program changes it to "leaker." But you really meant "speaker," so you press BACKSPACE and type it correctly. The next time you type "peaker," As-U-Type changes the typo to "speaker."


The AutoExpand function is a logical follow-up, letting you define macros to be expanded as you type. For example, when you type "thpb," the

program can replace that with "thrombophlebitis." This is convenient when pasting lengthy boilerplate text into e-mail messages and other correspondence. Macros can include the current date, time, and special characters, such as arrow keys.

As-U-Type also includes 10 distinct clipboards, but unlike Windows' standard clipboard, these are limited to storing only text. The QuickPad feature is handy for filling out online forms. It displays boilerplate text you define, pasting that text into the appropriate spot in another application with a double-click.

As-U-Type isn't infallible. In some circumstances, you can fool it into missing a misspelled word, so keep your old spell checker around for important documents. However, if your typing or spelling abilities are less than perfect, this little utility can squash embarrassing errors before anyone else sees them.

The program works with all versions of Windows since 95, but it isn't compatible with DOS applications or Microsoft Word's AutoCorrect Wizard. ▲



As-U-Type

\$29.95

Fanix Software

www.asutype.com





FrontPage 2002

\$169

Microsoft

www.microsoft.com



Microsoft FrontPage 2002

FrontPage is a solid WYSIWYG HTML editor. What's not to like? Nothing, as long as you design Web pages Microsoft's way, manage them on Microsoft servers, and publish them on servers supporting FrontPage extensions. Break free of this and you'll have continual headaches.

Creating a basic site from scratch without HTML knowledge is fast, although knowing some basics helps. FrontPage excels at adapting familiar elements from Office apps, and the toolbar, Help system, and layout will be familiar to Office users. In addition, documents, images, and files created in Office apps easily import to Web pages.

Microsoft has also improved the quality of the HTML code FrontPage generates. There's still significant code bloat in complex pages that FrontPage generates compared with Dreamweaver or GoLive, however.

Another notable FrontPage feature is the Browser Capability setting, which lets you target a specific browser when designing a page. Collaboration is FrontPage's weakest area; if you need to manage a site with multiple authors contributing code or graphics, consider exploring alternatives. Notably lacking is any integration of PHP or JSP. Overall, FrontPage is a solid, affordable choice for casual Web users. ▲



Dreamweaver 4

\$299

Macromedia

www.macromedia.com




Macromedia Dreamweaver 4

Dreamweaver combines robust HTML coding with a solid, configurable interface, plus seamless integration with Flash and Fireworks. Toss in hundreds of free Dreamweaver Extensions that leverage the work of thousands of designers, and it's easy to see why Dreamweaver has become the de facto Web site editor.

The app lets you simultaneously view source code and design layout; source code editing tools rival text-based HTML coding apps. Dreamweaver even includes tools to clean up HTML code that Word and other editors generate, and Target Browser Check generates a compliance report against specific browser types and versions.

Dreamweaver separates itself from the competition in site management. If your site is a team project, Design Notes, Site Reporting, and Asset Management tools make it easier to coordinate and maintain. If you haven't used Photoshop-style floating toolbars, Dreamweaver's interface can be daunting, but tasks follow consistent rules and you can modify the interface to keep toolsets on hand. The program's support for the WebDAV protocol lets you publish to supporting Web servers and use FTP. Dreamweaver's ability to create Flash objects, plus fluid integration with all of Macromedia's programming tools, makes it the reigning champ. ▲



GoLive 6.0

\$399

Adobe

www.adobe.com



Adobe GoLive 6.0

Ask a graphics pro to give up Photoshop and you'll have a fight on your hands. This is the type of loyalty Adobe wants for GoLive. With 6.0, the company just might earn it.

Adobe leverages Photoshop's interface with GoLive to seriously compete with Dreamweaver. Learning GoLive's nuances was easy, and the HTML code it generated was clean and compact. GoLive also handles Photoshop files in their native forms. Quick manipulation of DHTML and scripting languages is a must have for professional Web designers, not a bonus.

GoLive's support for authoring QuickTime 5.0 movies is sure to attract interest. Perhaps the

most compelling difference between GoLive and its competitors is the inclusion of Adobe's Web Workgroup Server. Utilizing WebDAV and SOAP, the Web Workgroup Server lets teams have a central repository for site management, versioning, and code maintenance. Any app (including other HTML editors) can access the contents of the repository because WWS is WebDAV-enabled. GoLive also supports WML and iMode devices. As 3G devices begin to penetrate the market, designing for these displays will become as crucial as designing for PC-driven browsers. GoLive is more expensive but presents a serious challenge for Macromedia. ▲

by Chris Pirillo

Snow White & The Seven Formats (+ 1)



Chris Pirillo is a geek in every sense of the word. When he's not distributing technology tidbits and eBooks on Lockergnome.com, he's hosting "Call for Help" on TechTV(callforhelp.tv). Every Sunday, he hosts a radio call-in program, assisting people with their technology troubles (Gnomeradio.com). This guy also has time to be a motivational speaker and drinks coffee to slow down. When he isn't preparing for his annual tech conference (Gnomedex.com), he's optimizing Windows XP's command line FTP tool.

I got a wild hair up my nose the other day and decided to start publishing eBooks. It all started when I moved from Iowa to San Francisco last year. I was lost. The first night in town I learned that calling the city "Frisco" was inappropriate.

"If only I had a little guide to tell me. And if only that little guide was digital."

Months later, I got tired of waiting for someone else to do it. I started to work on one myself. I took existing content from my daily newsletters and wrapped it into a nice 50-tip electronic document. After all, I didn't have a lot of time to research the Bay Area last week.

Collating the content was simple. Deciding on the final format was trickier. Which one would be universal? Which one would attract the largest audience? Which was the best? The answer: none. Therefore, I aimed to use them all. I began in Microsoft Word and hoped to locate affordable tools to take me everywhere I wanted to be (from there). Write once, read anywhere. I shall call this project Cappuccino.

And so began my journey. I gave each tip a separate page and a title header. The Table of Contents was generated on the fly. Easy.

Now, DOC works well, but it can leave some users a little grumpy (and I don't want to look dopey). Adobe's specific eBook product wasn't in the running because few people have it. I turned, instead, to a Microsoft Word plug-in that comes with Adobe Acrobat 5.0. Point, click, a flawless conversion to PDF. This brought me one step closer to my goal of providing the same content in the most popular digital formats. PDF is, for the most part, cross-platform. Bam! But I was far from finished.

Why not export the DOC as HTML? Sure, it'll take some cleaning and reformatting, but with a good template and the free Office 2000 HTML Filter (a source code cleaner available at office.microsoft.com), it will be a scalable operation. After I prepped it, I opened it in Internet Explorer and saved it as an MHT (with local image references). Check it out:

On any Web page click File, Save As, Save As Type, and Web Archive. Single file. Cool, huh? Yeah, but I wasn't finished with this HTML document yet. Microsoft also makes the HTML Help Workshop, which compiles CHM (Windows Help) files from Web pages. It didn't take long before I had two more formats to share with the world.

Exporting a Word document into text is a no-brainer; that takes care of the straight ASCII fans. With a few additional tags, I could have a TomeRaider document (viewable on Palm, Pocket PC, and Windows platforms). Simply add "<new>" before each new page in the text document and the TomeRaider software (shareware, tomeraider.com) pretty much takes care of the rest. What about Palmheads? PDB is the favored format, and there just so happens to be an abandoned piece of software floating out there that will create them for you

from straight text. If you can get a hold of BigDoc, it's a wonderful tool. If you can't find it, Express Reader DX is a great alternative (\$20; qvadis.com). To test the PDB, I downloaded EasyDOC Viewer (freeware, pilotid.com)

It's no secret I love my Pocket PC, which of course comes with the Microsoft

Reader built in. Heck, I even have the Microsoft Reader software installed on my desktop and laptop computers (even though I never use them on those systems). Amazingly enough, Microsoft also offers a Read in Microsoft Reader 1.1 Word document conversion utility (a quick search on Google will take you right to the page). In a few seconds, a Reader file will be sitting in your My Library folder, waiting to be synced and subsequently read.

DOC, PDF, MHT, CHM, TR, TXT, PDB, and LIT. If this "industry" really wants the eBook (eDocument, ePublication, eWhatever) to succeed, it better standardize on a single file format. Until that happens, I'll do my best to cover all of the bases. Oy! ■

You can dialogue with Chris at chris@cpumag.com.

And so began my journey. . . . The Table of Contents was generated on the fly. Easy.

The Nature Of Computer Security:

An Interdisciplinary Approach



Pete Loshin, former technical editor of software reviews for BYTE Magazine (print version), consults and writes about computing and the Internet. He also runs www.linuxcookbook.com. He owns shares of both Microsoft and Red Hat and believes that Windows isn't for everyone, but neither is Linux.

Want a practical, 100% secure computer or network? Sorry, it's a mathematical impossibility, open source or proprietary. Yet, reliable security is often cited as a reason to switch from Windows to Linux or BSD. What gives?

Most of us think of passwords, firewalls, viruses, encryption, and digital signatures when we hear "computer/network security." But technologies and products don't guarantee security. Anyone claiming his product makes systems uncrackable is either deluded or lying.

Amazon.com lists over 1,500 titles under "computer security." Unfortunately, few are written by bona fide security experts (one that is: "Network Security: PRIVATE Communication in a PUBLIC World," Charlie Kaufman, Radia Perlman, and Mike Speciner). Yet the best security books aren't always shelved with the others.

"Godel, Escher, Bach: An Eternal Golden Braid," by Douglas R. Hofstadter repays the diligent reader seeking to understand the nature of computing. Hofstadter explained Godel's Theorem and showed me that unbreakable complex formal systems are impossible.

Godel's Theorem states that no complex formal system—a framework for systematizing mathematics, a programming language, a computer—can be proven "complete," that is, capable of defining all of its own possible inputs, outputs, and states, and excluding the one that breaks the system.

If, as Hofstadter suggests, you consider a record player as a formal system, it is complete if it can reproduce any possible sound. There will always be at least one sound that vibrates the thing to bits. Stop the vibrations at one frequency, and because the wood, plastic, metal, and glass components are vulnerable to vibrations, there will always be some other frequency that can break your turntable. If your record player is made from atoms, it is vulnerable to vibrations.

Computers are complex formal systems: They accept input, produce output, and have many states—some of which will crash the system. The attacker uses input (virus, buffer overflow, trojan) to put the system in vulnerable states.

Any program or computer must accept inputs and produce outputs based on internal state to be useful. But by accepting inputs, you must also accept that one or more of those inputs will break your system. Patch the system and you merely trade known killer inputs for unknown ones.

Cooties

We've been doing computer security for a few decades, but organisms have been defending themselves from attack for eons. Evolution and epidemiology demonstrate the importance of good defense. In "Guns, Germs, and Steel: The Fates of Human Societies," Jared Diamond shows how disease and the ability to resist it shape history. Understanding how organisms defend themselves sheds light on computer defense.

For example, the Ebola virus reproduces rapidly, spreads easily, and kills up to 100% of those infected. Most outbreaks occur in remote areas, but since Ebola incapacitates infectious victims swiftly, outbreaks "burn out" quickly before they can spread.

A virus, whether computer and biological, must find an appropriate host, infect it, and reproduce. A computer virus transmitting one copy before destroying its host would burn out as soon as it attempts to infect an immune host. A virus that tries infecting 100 targets can spread faster and farther than less virulent virii.

Computer virii are platform-specific, Windows being the platform of choice, so Windows-based virii can spread faster and farther. With 90% of all desktops, on average, a new virus can infect on 90 out of 100 systems it attempts. By the fourth generation, as many as 25 million systems could be infected.

For now, *NIX-based virii could infect only about two of 100 desktops. The first infected host infects two others; they tell two friends, and they tell two friends, and in only 10 generations you've infected about 1,000 other hosts.

Effective predators (yes, a virus is a predator) can ravage homogenous populations. Crop pests, from potato borers to (cotton) boll weevils, abound on farms because that is where the potatoes and cotton are located. A virus lethal to one farm-raised chicken will be lethal to millions of chickens.

With relatively little genetic diversity, one farm-raised chicken is much like the next chicken.

Diverse networks can withstand a virus that would devastate single OS networks. Which brings us back to open source: Add diversity with Linux, BSD, or commercial versions and you improve your network's ability to withstand an attack. A Windows virus won't bother Linux users, and a BSD virus won't affect Windows users. Cyber-diversity confers a degree of immunity from otherwise catastrophic attacks.

Killer Cooties

Godel proved that systems sufficiently complex to be useful are implicitly insecure. Computing systems accept input and produce output. That some input will "break" any formal system means you can plug the hole only by completely forbidding all inputs—rendering the system useless.

But what if you tried to block harmful vibrations from your record player instead of changing its components or design? The attacker wouldn't be able to penetrate a good firewall-style barrier, right? Wrong. Your barrier is vulnerable; vibrate the barrier first, then on to the record player. You make a bigger shield, they make more noise.

This is an arms race, which is a costly affair. If GEB:EGB makes your brain hurt, read Dr. Seuss' "The Butter Battle Book" for a look at arms races even grownups can understand.

Organisms have arms races, too. A species that has no natural predators

eventually loses its defenses (why develop fast legs if it never runs from anything?). Conflict can speed things up, so a prey species either develops defenses or it will get eaten.

So, do you monitor your perimeter closely and leave things as they are, or lock down every new hole as it's discovered? You can lay traps by the obvious mousehole in the kitchen, or seal it and force the mice to look for a trickier (and harder to detect) way in. Mice or crackers, the ones that get in reproduce by making more mice or by writing and posting attack scripts.

Laurie Garrett's bloodcurdling book, "The Coming Plague: Newly Emerging Diseases in a World Out of Balance," shows what happens when you engage in an arms race with a virus. Inappropriate antibiotic use is winnowing out weak bugs and turbo-charging the evolution of super-cooties.

**So, do you monitor
your perimeter closely
and leave things as
they are, or lock down
every new hole as
it's discovered?**

Killer rabbits might evolve in an environment filled with prey but short on rabbit food. We could breed a killer rabbit, too: put a bunch in a closed environment, introduce increasingly effective predators, and weed out the weak ones. Each predator kills most of the rabbits leaving only the biggest, toughest, and roughest rabbits to breed. Cats first, then dogs, jackals, hyenas, wolves, cheetahs, lions, grizzly bears, alligators, tigers, giant squid, etc. Your new rabbit will eat sharks for breakfast.

Scared? Well you should be. Our antibiotics kill off the wimpy bacteria and leave only the toughest ones with special defenses. So if you catch a drug-resistant strain of TB or anthrax, that's all, folks.

Computer viruses can't participate in their own evolution as real virii do (by mutating), but the systems they target certainly can be adapted to block virus threats. If you're a biological bug, the drug is the enemy: evolve or die. Systems have the same option: evolve or become obsolete. *NIX OSes have evolved over three decades (that's 12 billion years in Internet time), plenty of time for attackers to scrutinize the code for vulnerabilities and defenders to plug them.

Perfect vs. Good Enough

In security, perfection isn't only wasteful, it's impossible. Encryption algorithms all fall to brute-force attacks with sufficient force. As faster hardware makes those attacks more practical, defenders rely on longer keys. Even if perfection on one aspect of security were possible, it would focus attackers on other aspects, to seek out weaknesses. It's much easier to bribe a disgruntled employee than it is to crack a strong encryption.

As a practical matter, open source OSes offer security advantages for many reasons not always entirely attributable to superior technology. Can Windows ever be more secure than OpenBSD? Can Linux ever be more secure than Windows? Can any OS ever be "more secure" than another OS? None of those questions matter. Just this: Is your OS security good enough for how you use it?

Knowledge is a wondrous thing, and you don't always find it where you are looking for it. Let me know if you discover more examples where interdisciplinary cross-pollination is shedding light on technology.

(Pete Loshin's lawn is a model of biodiversity, supporting at least 38 distinct species of plant. Sadly, all of those species are weeds.) ■

You can get saucy with Pete at opensauce@cpumag.com.

Words From The Web

Yeah, they *actually* said this . . .

From an MSN celebrity chat with Jack Black, upon being asked how he balances his acting and singing careers:

"It all gets mooshed into one entertainment burrito."

With extra cheese.

From a Jolt.com message board about jobs:

"Burger king sux. S*ty pay and at the end you smell like onions. And if you're caught high YOU GET FIRED!"**

Apparently you can't always have it your way, even at Burger King!

From a Maximum PC message board:

CPU is a totally kick a magazine!!**

That's right, baby. That's right.

From a Yahoo! advice chat room:

"My boyfriend is the one who's a little nutty, but I love him."

Who is nuttier? The nut, or the woman who loves him?

From a Malgoo.com message board on hobbies:

"I use to have an Atari 800XL with games on tapes, you use to have to wait hours for the games to load and when they did they never use to work anyway."

But look how far Atari's come since then!

Livin' Large In Ladonia

Ever feel like you need to get away? Try punching a one-way ticket to Ladonia. Ladonia hit the news wires in March 2002 when it was reported that many Pakistanis wanted to immigrate there, only to discover that Ladonia doesn't exist. It's a mostly online-dreamland (www.ladonia.net) created by Swedish-seceding artist Lars Vilks in 1996. Vilks staked out a 1-kilometer (0.62 hundredths of a mile) swatch of Swedish countryside and declared it to be a sovereign nation.

Several people have become citizens of Ladonia since then, and you can too. All you have to do is fill out an online form. The most stringent requirement is that you provide a Latin word or phrase when you fill out the form to help build the language of Ladonia. You can also buy your way in for \$12.

Some points of interest in Ladonia: The national

anthem is the sound of a stone being thrown into water, one Örtug is equal to about 10 Swedish crowns, and the Ladonian flag is green with a green cross. Get your passport today. ▲

My friend went to
LADONIA
and all I got was this
T-shirt.



Infinite Loop

My Life's Work

After Bill Gates' famous "security" e-mail message leaked out Jan. 15, rumor had it that Microsoft's employees would spend the entire month of February combing the company's products for security flaws.

Say all Microsoft employees searched for bugs (highly unlikely) each of the 20 business days in February, assuming eight-hour workdays (also highly unlikely). Microsoft reported a 48,958 headcount at the end of January. The old Microsoft Calculator puts the figure at about 7,833,280 man-hours, or approximately 894.21 years total. That's almost 12 average persons' life-times spent just searching for flaws in Microsoft's products.



This Site's Done Right

OK, so Derek's Web site (below) may not win any awards for innovation, but here's one that has won several awards. The Templar Studio Web site (www.templar.com) has one of the coolest home pages I've seen in a long time. Templar Studios creates Web-based interactive entertainment (read: games), primarily for Lego toys.

The home page is an animated office building that contains bios of Templar Studios staff members, as well as a few hidden features. Move your mouse pointer over one of the staff members, and he or she will spring into action. You can get the staff member's bio by clicking his or her image.

The staff bios are well done, but the other information on the site is pretty cool, too. There

are several objects throughout the building that come to life when you move your mouse pointer over them, such as a TV screen, a jukebox, and a phone. Click the TV screen to see a few animation clips developed by Templar Studios. I liked the "Stormrunner: Dark Planet" clip best. The Jukebox area contains several audio clips that Templar created for various endeavors. Click the phone on the desk of programmer Nate Acheson (first floor) to find Templar's contact information. You can also scroll through a gallery of Templar Studios' work by clicking the framed pictures on the wall in the first-floor office. This crew has done some pretty interesting stuff, not the least of which is the company's Web site. ▲



Stuck In Wally World

We know there are some pretty boring Web sites out there: family Web pages filled with pictures of puppies and Ford Festivas, the occasional blowhard treatise on the emotional import of Henry Rollins poetry, new computer magazine Web sites that don't have a lick of content. But some Web sites are so pointless, so bland, so inconceivably void of purpose, you're compelled to check them out. That's how I felt when I came across Derek's Big Website of Wal-Mart Purchase

Receipts (lightning.prohosting.com/~receipts/index.shtml).

The Web site consists almost entirely of scanned images of receipts, one of the most pedestrian, ubiquitous objects known to modern man. The receipts are accessible through a master index that lists the date of each, from late 1996 to 2002. Visitors to the site can post comments for each receipt, which they do in abundance. Of course, Wal-Mart is not officially associated with Derek, except that Wal-Mart is his main supplier of Mountain Dew Code Red. Derek explains in his disclaimer why he chose Wal-Mart's receipts: "This site is not sponsored, ordained, sanctioned, or provided by Wal-Mart. Really, its only relation to Wal-Mart is that Wal-Mart prints out the clearest receipts. . . . If my grocery store's receipts didn't fade so bad, this would reference Hornbacher's instead."

I scrutinized every last receipt, and in doing so, I think I've learned a lot about Derek and Wal-Mart. To wit:

- Derek almost always buys boxes of facial tissue in threes.

- Between April 1, 2000, and June 13, 2000, the price of facial tissue skyrocketed from 84 cents to 88 cents.
- Derek is a fine gardener, or at least he buys plenty of gardening supplies so that people will think he is.
- Apparently, Wal-Mart sells a two-story condo for \$19.97. (Or so it was listed on one receipt. That's value!)
- Derek's really into photography, as evidenced by the many rolls of film he buys and has developed at Wal-Mart.
- Wal-Mart changed the design of its receipts between June 18, 2000, and June 29, 2000. Frankly, the new designs are better.
- Greeting cards are listed as Counter Cards on receipts. What kind of freaky Wal-Mart-speak is that?
- Derek has way too much time on his hands.

Derek's Big Website of Wal-Mart Purchase Receipts also has a guest book and a pretty funny FAQ, along with links to pro-Wal-Mart and anti-Wal-Mart sites. ▲

If you find a strange, interesting, or funny Web site in the course of your Internet travels that you think is worthy of Fringe, send your suggestion to fringe@cpumag.com



TIME WASTERS

THE BATMOBILE IS JET-POWERED.
BULLETPROOF. FIREPROOF.
AND REACHES SPEEDS OVER 300 MPH.
COULD WE POSSIBLY IMPROVE ON IT?



YES.



NOW THE VEHICLE THAT HAS EVERYTHING HAS ONSTAR®. IF BATMAN'S AIR BAG DEPLOYS, WE'LL RECEIVE A SIGNAL AND CHECK ON HIM. IF A VILLAIN STEALS THE BATMOBILE, WE'LL TRACK IT RIGHT AWAY AND NOTIFY THE POLICE. WE CAN EVEN SEND REMOTE SIGNALS TO UNLOCK HIS DOORS OR DIAGNOSE HIS ENGINE WHILE HE'S DRIVING. YOU CAN'T GET YOUR OWN BATMOBILE. BUT YOU CAN GET ONSTAR ON YOUR NEXT VEHICLE. TO FIND OUT HOW, CALL 1-888-ONSTAR-7 OR VISIT WWW.ONSTAR.COM.



HOW CAN WE HELP YOU?

BATMAN, THE BATMOBILE AND ALL RELATED CHARACTERS, NAMES AND INDICIA ARE TRADEMARKS OF DC COMICS ©2001. ©2001 ONSTAR® ALL RIGHTS RESERVED.



Searching The World, Finding Success

This was the problem: "Your search for 'elephants' + 'Poughkeepsie' generated 6,497,152 hits."

Until the late 1990s, searching the Web was an arduous, time-consuming affair dominated by page after page of irrelevant returns. Early search engines were easily duped by frequent word use. Simply fill the bottom of your page with hundreds of lines of key words made "invisible" by turning the text the same color as the background. Dumb search engines figured that with 20 instances of the keyword appearing in the page, it must be a relevant hit. Then the same tactic was more or less applied to meta tags within the HTML code, which is why most search engines now ignore meta tags that might otherwise be put to good use.

The situation seemed hopeless as search engines seemed unable to find an effective method for coping with the chaotic, sprawling landscape of online data. Many Web users were starting to turn sour on the prospect of using the Internet as an information tool. Then in September 1998, a beta search engine named Google appeared, and everything changed.

Google Roots

"Google" is a twisting of the word "googol," a term coined by American mathematician Edward Kasner's young nephew Milton Sirota to mean the number 10^{100} , or 1 followed by 100 zeros. (Impress your friends at the next trivia match by knowing that the word for 10 to the power of googol is a "googolplex," not to be confused with "the Googleplex," Google's Mountain View, Calif., headquarters.) Although this might have been a sly, hyperbolic jab at the number of false hits most search engines were returning, Google's founders, Sergey Brin and Larry

Page, use the term as an allusion to the unfathomably large amount of online data.

Brin and Page started researching search technology in 1995 as part of their doctoral research at Stanford University. According to a paper the two wrote titled "The Anatomy of a Large Scale Hypertextual Search Engine," one of the foremost search engines in 1994 was the World Wide Web Worm (WWW), which indexed 110,000 Web pages and documents and received roughly 1,500 queries per day. It was clear to everyone that these numbers were due to grow exponentially but that contemporary indexing and searching systems were inefficient at best. In fact, the paper notes, when their paper was written in November 1997, only one of the top four commercial search engines listed its own search page within the top 10 hits when searching for itself.

When their collaboration began in earnest in early 1996, Brin and Page, then 23 and 24, nicknamed their would-be search engine BackRub and worked on it while pursuing their degrees. By the summer of 1998, they had to make a choice. It was time to get serious about bringing BackRub to the world, but the time commitment involved would swamp their scholastic efforts. Encouraged by Yahoo! co-founder and fellow Stanford alum David Filo, plus 1 million dollars in private investment backing, Brin and Page put their Ph.D.s on hold and set their newly renamed project, Google, loose on the public.

Google's Heart

One year after the beta launch, Google had attracted another \$25

million in private funding. Almost immediately, Google became a favorite destination for search-savvy 'Netizens, although it remained something of an underground phenomenon, trailing in the shadow of search giants such as AltaVista, HotBot, InfoSeek, and Lycos. But what distinguished Google was the uncanny accuracy of its hits.

PageRank. Google's search methodology hinges on a patented technique called PageRank, which attempts to judge the quality of a page in part by the number of other pages linking to it. If page X links to page Y, Google sees this as one "vote" for Y. More votes equate to higher quality pages. The supposition is that no one would bother linking to a page unless it contained something of merit.

Complementing this, Google also examines the page casting the "vote." A link to a URL from a home page on a site such as Yahoo! or MSNBC is almost sure to guarantee top ranking for that URL. This helps to minimize manipulation of the system by those who would simply create hundreds of bogus pages linking to a single site. Another



advantage of analyzing links (or, as Google often calls them, anchor text) is that the engine can index and rank items and pages without any text, such as databases and images. Engines that rely exclusively on keyword frequency, for example, would entirely miss a page containing only images, even from a high-profile artist such as Ansel Adams or Picasso. By indexing link targets, Google can return hits for pages it hasn't even directly examined yet.

According to Brin and Page, an average workstation can compute PageRanks for 26 million Web pages in only a few hours. However, simply assessing a page's quality doesn't translate into search relevance. The page's text still needs to be analyzed.

Association. Keyword association now comes into play, but Google does more than just match character strings. The search engine also inspects font attributes such as size, boldface, and italics. Words with such characteristics, obviously meant to emphasize the text, are given more weight. Google also takes subhead divisions into account as well as the location of a word on the page. The engine even

takes into account the content of neighboring pages.

In the beginning, Google asserted that its relatively small catalog of Web pages was compensated by the accuracy of its results. Today, Google need make no excuses. The service indexes more than 2 billion unique pages, the largest of any engine on the Web, and processes more than 150 million search queries per day.

Using Google

Google searches appear to be much like most other Web searches, but there are some key differences. Sure, if you go to www.google.com, type in a search phrase, and press ENTER, the engine will return a page full of hits. But whereas many engines assume that you are searching for any of the words, Google searches for all of them. In Boolean search terms, every Google search assumes there is an AND command between each word. However, you can use an "any" approach by inserting OR (in capital letters) between appropriate terms. For example, if you wanted information about ancient Central American ruins, you might

do a search for "ruins Aztec OR Mayan." Also note that, unlike most engines, Google is not case sensitive with search words.


Another useful Google feature is the ability to include stop words. A stop word is a character string that appears so frequently that it is essentially useless to search for it, so the engine simply ignores it. Different engines eliminate different stop words, but it's a safe bet that if you search for "1 night with me and Scooter," the only words the engine will process are "night" and "Scooter." All other terms are too common to be useful.

That said, if you're sure there is a page out there with those specific terms or that phrase, you can force Google to recognize stop words with the + symbol. So you would type "+1 night +with +me +and Scooter" (without quotation marks). To search for the entire phrase, you would use quotation marks around the entire phrase without the + symbols. And don't worry if you spell it "nite" instead of "night." Google integrates spell checking.

Special searches. Try the Google home page's Advanced Search link for even more

Top Brass: Sergey Brin, co-founder and president of technology for Google

CPU: Why did you pursue developing a search engine rather than something else?



Brin: We believed that search was a very important application and becomes more important as the Web continues to grow. Our research proved that our technology returned more relevant, high-quality results than any other existing search engine at that time. We continued to develop our

search technology and eventually created PageRank technology, which is the heart of Google today.

CPU: How tough was it to keep Google's home page ad- and clutter-free when the market trend was to do exactly the opposite?

Brin: It wasn't tough at all. Google's ad philosophy is built on highly targeted ads only, which means that users only see ads that are relevant to their search. Placing ads on the home page is ineffective because we don't know what a user is intending to search. In addition, Google's simple,

clean, and uncluttered design enables users to search the Web easily and quickly with no distractions.

CPU: When did you know Google had "turned the corner" and was going to make it to the big time?

Brin: Google has not yet made it to the big time. We will know once we accomplish our mission to organize the world's information and make it universally accessible and useful.

CPU: Look into your crystal ball and tell us what Google will be doing three years from now.

Brin: Google is constantly innovating to provide access to more information in a multitude of languages and formats. We will continue to index the most useful Web pages, as well as other types of information (such as images, PDF files, etc.), as the Internet continues to grow. Google will continue to explore opportunities on numerous wireless platforms. Location-based technology is another area we are currently researching. As we stay true to our mission, we will continue to develop advanced technology that enables users to search the entire Web easily, quickly, and effectively. ▲

options. If your query deals with the U.S. government, Apple, Linux, certain universities, or other niche areas, try one of the Google engines specifically devoted to that area to boost the relevance of your hits. Use the date range function if you only want hits from a certain time period, a handy tool if you want only current events pertaining to a certain person, for example. If you have kids using Google, consider using the SafeSearch filter to weed out hits with potentially objectionable content.

Be sure not to miss Google's peripheral databases. The Image Search archive currently indexes more than 330 million

is no longer active, you can access a prior version of the page in Google's archive.


Get Googlified

According to a recent MSNBC article, Google has passed the true test of American pop culture by passing from being a proper noun into a verb. People "google" one another to confirm their identities before a blind date or engage in the game of "Googlehacking," or trying to find a two-word search that only returns a single hit.

Perhaps soon, another application of the company's name will be as a benchmark for online integrity. Google is famous for

Ph.D.s and roughly 80 engineers. Currently, the company is focusing on expanding its business on two fronts. The first is to gain a wider international following, which includes opening offices in Europe and Asia. The second is to become the premier search engine in the burgeoning wireless space. Google already has deals with Palm, Handspring, and Japan's DoCoMo and will no doubt be at the forefront of smartphone-based searching as that technology grows over the next few years. Look for speech recognition to be a main ingredient in Google's future wireless success.

The Life Of A Query



SOURCE: GOOGLE INC.

pictures, again making the largest such catalog on the Internet. Another Google size prizewinner is its Usenet archive. Google Groups archives more than 700 million newsgroup messages, spanning thousands of topics and dating clear back to 1981. Research freaks or the socially curious would do well not to miss this archive, which predates the Web by more than a decade and contains human commentary on practically every subject imaginable.

There's more: The engine archives more than 35 million non-HTML files, such as DOCs and PDFs. You can filter hits to only display results in any of 35 languages, as well as translate international pages into English. Google's Street Maps will reward your search for "1600 Pennsylvania Ave. Washington, D.C." with a street map of the White House area. Another handy feature is that Google caches HTML content, so even if a result link is inactive because the destination's server is down or the page

its ultra-plain home page and dearth of obnoxious ad content, a strategy that undoubtedly cost the company money in the short term but earned the respect and trust of millions of Web surfers. Unlike other engines, Google does not offer higher rankings to advertisers, instead deriving the bulk of its revenue by licensing its search technology.

Despite many rumors to the contrary, Google remains a privately held company. Sergey Brin is credited in a June 2001 BBC News story as saying that an IPO would offer Google "additional respect, more cash, and stock liquidity." Nevertheless, the two founders and new CEO Eric Schmidt have kept Google private, both because the company still has plenty of cash left from its initial funding rounds and to keep the company focused on marketing and research rather than keeping Wall Street happy.

Research and development are key in this company that sports more than 30

Naturally, Google's future won't be all sweetness and light. The search engine is gaining enough popularity that major portals such as MSN, AOL, and current Google client Yahoo! may view Google as more of a competitor than a complementary service. Losing such present and potential clients may have a severe effect on Google's bottom line. Similarly, while Google has been in a searching league of its own for the last three years in terms of quality hits, competitors such as Teoma and SearchOnline are racing to catch up. As Microsoft has demonstrated repeatedly, long-term success may have less to do with product quality than partnerships and marketing savvy, areas in which Google is still learning the ropes. If Eric Schmidt can keep Google on the path to searching sovereignty, though, there seems little that the company won't be able to accomplish. **CPU**

by William Van Winkle

Defining XML Dialects

In *Coder's Corner: XML*, Ian Graham shows you how to program with XML. Ian is the author of numerous books pertaining to Web development, including *"The HTML Sourcebook"* and *"The XML Specification Guide."*

Previously, we've introduced XML's main concepts—including syntax and the XML name space mechanism—and examples of XML dialects. Now we'll look at defining specific dialects with Document Type Definition, or DTD, which is defined in the XML 1.0 specification. Until recently, this was the only way to specify such rules. The W3C (www.w3c.org) has released a new specification, XML Schemas, designed to eventually supplant DTDs. We'll save Schemas for another time.

Rules for writing DTDs are complicated. We can't cover them all here, so we'll focus on basic mechanisms. A DTD essentially lets you define:

- The names of the elements, such as tag names, allowed in a document and the structural relationship of those elements. For example, a mailServer element must come inside a configPar element, not before it.
- The names of the attributes each element supports, plus rules restricting the attributes' values.
- The entities, or special quantities representing blocks of text and/or markup, you can use to insert these blocks into an XML document (similar to a macro expansion facility).

In a document, such rules are defined inside a special markup block called a Document Type Declaration, which also has the DTD abbreviation. This declaration must always appear at the beginning of the document (after the XML declaration), as in this example:

```
<?xml version="1.0" encoding="utf-8" ?>
<!DOCTYPE configPar [
  <!ENTITY contentStuff "<mailServer>
pop3.serversite.org </mailServer>" >
```

```
]>
<configPar>
  &contentStuff;
</configPar>
```

The declaration begins with the special string, `<!DOCTYPE`, followed by the document type this describes (`configPar`). A left bracket follows and then the DTD rules. The Document Type Declaration ends with the sequence, or `]>`, which appears before the start tag. Note that the document's root element, `configPar`, is the same name as for the start of the declaration. If these aren't the same, there's an error in your declaration.

This example DTD uses an `<!ENTITY` declaration to define an entity named `contentStuff`, with the content of the entity being the string in the quotation marks. This declaration type (there are other entity types) defines an entity you can use inside the XML document, using the notation `&contentStuff`. This example has such a reference inside the `configPar` element. When an XML processor reads and processes this document, the data sent to the program that reads in and uses the data will correspond to:

```
<configPar>
  <mailServer> pop3.serversite.org
</mailServer>
</configPar>
```

The entity will be replaced by its content. You don't have to set the entity's value inside the DTD—an entity declaration can instead reference (using a URL) an external resource containing the data. For example, suppose you're writing a book (using XML) and have the preface and various chapters stored in files, such as `preface.xml`, `chap1.xml`, and so on. You could write a XML document, such as `book.xml`, that aggregates the components using external entity references. Here's an example:


```
<!DOCTYPE book [
  <ENTITY pref SYSTEM
    ".preface.xml" >
  <ENTITY chap1 SYSTEM
    ".chapter1.xml" >
  <ENTITY chap2 SYS
    TEM "http://www
      .iangraham.com
      /book3/chapter2.xml" >
]>
<book>
  &pref;
  &chap1; &chap2; ... and so
  on ...
</book>
```

The entities `pref` and `chap1` refer to files adjacent to the `book.xml` document, while `chap2` references a resource at the indicated URL. The XML processor processing the `book.xml` file will retrieve these resources and replace the entity references with the data in them.

I Declare

As mentioned, DTDs can contain declarations that define rules for placing elements in a document, and for the attributes those elements can have. Here are examples:

```
<ELEMENT configPar (mailServer+)>
<ELEMENT mailServer (#PCDATA
  TA)* >
  <!ATTLIST mailServer
    type (pop3 | imap3 | imap4)
    "imap4"
    webMail CDATA #IMPLIED
  >
```

Element placement rules, or content models, are defined using `<ELEMENT` declarations. The first line states that the `configPar` element must contain one or more `mailServer` elements (the plus sign `[+]` denotes this). Conversely, the `mailServer` declaration states that the `mailServer` element can only contain text (the keyword `#PCDATA` means regular text); it can't contain markup tags.

The `<ATTLIST` declaration defines the attributes supported by the `mailServer` element. The example defines two attributes: the attribute `type`, which can have the values `pop3`, `imap3`, or `imap4`, with `imap4`

being the default, and the attribute `webMail`, which can have any string (CDATA) as its value. The keyword `#IMPLIED` means this latter attribute is optional and you can omit it.

This part of the DTD can add attribute information. For example, the attribute declaration for `"type"` states that the default

data you can accurately turn into data by an XML parser. Being well-formed is the minimum condition for being XML.

Well-formed documents don't have to have a DTD, but they can. For example, a well-formed document can have a DTD that defines entities used inside the document. This is convenient for defining reusable markup components, such as entities representing common special characters or phrases. A DTD can also define default values for attributes.

A valid document has a DTD, has document grammar rules completely specified by the DTD, and is document consistent with those rules. Validity is a stronger constraint than being well-formed because it guarantees the document is well-formed and structured according to the DTD's rules.


Validation is useful when editing a specific XML dialect. The rules ensure that all tags are in the right places. However, you don't always need to enable validation. Web browsers, for example, simply check to make sure the data is well-formed and typically don't check for validity.

This practice has led to two models for XML parsers: validating and nonvalidating. A validating parser checks a document rigorously against the DTD. If there are inconsistencies, the parser flags an error and stops. A nonvalidating parser processes the DTD, looking for entities it may need (or attribute declarations) and processes the XML data, including the extra information the DTD provides. It won't check for grammatical consistency (other than for it being well-formed), however.

We haven't discussed XML names because DTDs don't support the XML name specification, and there isn't an easy way to include name information in a DTD. This was a reason for creating XML Schemas, which let you define grammar rules for XML documents, but also use name spaces in those definitions. **CPU**

by Ian Graham

(See the full examples of those in this article at www.smartcomputing.com/cpumag/jun02/coder or at www.utoronto.ca/ian/jun02)



XML Spy 4.3 (www.xmlspy.com) is one application that can give you an editable, graphical view of the rules specified in a DTD.

value is `"imap4"`, which means the attribute exists and has this value even if it's not explicitly set in the document. An XML processor, after reading this DTD ahead of our example document, effectively changes `mailServer` elements to start with `<mailServer type="imap4">`, adding the default value from the DTD's attribute declaration.

Often, including a DTD doesn't make sense. Instead, the Document Type Declaration can reference an external location for the DTD using a URL. Using this form, the `configPar` example becomes:

```
<?xml version="1.0" encoding="utf-8" ?>
<!DOCTYPE configPar SYSTEM
  "http://www.where.org/stuff/configPar
  .dtd" >
<configPar>
  &contentStuff;
</configPar>
```

The file `configPar.dtd` contains all the declarations formerly in the `DOCYTPE`. You can use this DTD again across all your documents. The app reading each document will retrieve and process the DTD.

Well-Formed vs. Valid XML

A syntactically correct XML document with no XML markup syntax errors is "well-formed." This formal term is defined in the XML specification to indicate XML

Web Design Workflow: Part II

In last month's issue, we discussed Part I of the Web design workflow: starting from scratch and building a "wish list" of content and then refining that list into a workable outline. Your outline, when finished, can then morph magically into a "sitemap"—a flowchart-like diagram of your Web site that shows its overall information structure.

Building A Sitemap

To build a sitemap, have your completed outline handy. Start by drawing a small box (like .5 inch x .5 inch) at the top of a document. (I suggest using an illustration program like Adobe Illustrator, Macromedia Freehand, or anything that allows you to build simple shapes and lines.) This small box at the top represents your home page, so add a label with the site's name such as "The Four Food Groups."

In a row below, add a series of small boxes that represent the main category headers from your outline. Remember, for usability reasons, your outline should have about five to seven main sections, so your row will include five to seven boxes in a horizontal formation. Put a label next to each one. Draw a line that goes from the home page box to each of these boxes in the second row.

Below each of these boxes, stack a vertical column of boxes representing each of the subsections for a category. For example, if "Vegetables" is a main category of your "Four Food Groups" Web site, make a series of boxes, stacked vertically under the Vegetable box with labels like Tomatoes, Carrots, and Peas. Basically, the idea is to show the outline's hierarchy in a series of rows with boxes nested under other boxes. Again, draw a line that connects the Vegetable box to each of the boxes stacked below it.

Building "Wireframes"

Although a sitemap is useful for giving you the 30,000-foot view of your Web site, it does not show

you the structure of each *page*. For that you need a set of "wireframes." A wireframe is a diagram that shows the contents of a Web page—from navigation elements to body copy and Flash movies.

Building wireframes is an important step in the Web design process because it forces you and the

client to itemize all of the *content*, like navigation bars, movies, and body copy, that will be needed for each page of a Web site. In a sense, a set of wireframes becomes a content "laundry list" and later, a blueprint for building the site itself.

Think about it, after building a complete set of wireframes for a Web site, suddenly you and the client have a document that shows what content

needs to be written or assembled for the whole site. The wireframes, therefore, underscore how critical the client's role is in the creation of the Web site—it's their job to provide all this content! Not only that, the wireframes also show the client how much space each bit of content occupies relative to the rest of the page. This gives them a better sense of how *much* content to write and assemble for each page.


In next month's article, I'll dig further into tools and tips for building effective wireframes and show how wireframes play an integral part of the Web production workflow that follows. ■

WARNING:

This article contains excerpts and concepts presented in "Web Design for Dummies." Don't let the name fool you, the book is designed for creative professionals, and there are a lot of advanced Web design topics and issues covered. Plus, it's a fun read. You can contact Lisa at lopuck@cpumag.com and see her work at www.lopuck.com.



Lisa Lopuck, www.lopuck.com, is a Web creative consultant helping companies to define and plan their Web creative strategy, information flow, and visual look and feel. She is also the author of numerous best-selling books on Web design, including "Web Design for Dummies," and is a sought-after speaker at Web conferences and universities around the world.



PROOF

THAT *great* IDEAS
GET HEARD.



HONORED BY *FORBES* ASAP for CHANGING THE WORLD.

Forbes ASAP magazine rates the **Bose® Wave® radio** right up there with the light bulb, the compact disc and the telephone. Only fifteen inventions made their exclusive "All-Time A-List" of technology breakthroughs that have changed the world, and the Bose Wave® radio is one of them.



Recognizing products "both functional and beautiful," *Forbes* ASAP credits Bose with managing to fit "really big sound into a really small – and pretty – package."



LIFELIKE SOUND

The secret is Bose acoustic waveguide speaker technology, which is


built into both the Wave® radio and Wave® radio/CD (shown above). This and other patented Bose technologies deliver rich, full sound with remarkable clarity and deep, natural bass. It's sound you'd never expect from a system that fits so neatly on a tabletop or kitchen counter.

HEAR the DIFFERENCE

The Wave® radio and Wave® radio/CD are available directly from Bose, the most respected name in sound. Both are available for **12 low interest-free monthly payments.***

Order before June 30, 2002 and receive a free Bose CD carrying case that holds 24 CDs. Our risk-free, 30-day, in-home trial guarantees

your satisfaction. You really have to hear one for yourself. So call today. Because the proof is in the listening.



FREE CD CASE
Order before June 30, 2002.

Call Today
1-800-836-6754
ext. T8941

For more information on all our products:
www.bose.com/t8941

Please specify when ordering.
Wave® radio/CD: ☐ Platinum White or ☐ Graphite Gray
Wave® radio: ☐ Platinum White or ☐ Graphite Gray

(Mr./Ms./Mrs.)
NAME ADDRESS
CITY/STATE/ZIP
DAY PHONE EVENING PHONE

Mail to: Bose Corporation, Dept. CDD-T8941
The Mountain, Framingham, MA 01701-9168

BOSE
Better sound through research®

©2002 Bose Corporation. Patent rights issued and/or pending. The Wave® radio design is also a registered trademark of Bose Corporation.
*Installment payment plan and free CD case offer not to be combined with any other offer or applied to previous purchases.
Payment plan available on credit card orders only and subject to change without notice. Quote: *Forbes* ASAP, 11/27/00.

netradio.txt



Rob "CmdrTaco" Malda is the creator and director of the popular News for Nerds Web site Slashdot.org. He spends his time fiddling with electronic gizmos, wandering the 'Net, watching anime, and trying to think of clever lies to put in his bio so that he seems cooler than he actually is.

It's time I come clean: I am addicted to music. I buy countless CDs. I download music from the 'Net using whatever P2P system functions under Linux. I have music playing whenever I'm using a computer in my office. My co-workers and I have a giant hard drive where we've put every CD we own. It just doesn't make sense to dedicate the shelf space to jewel boxes when a couple of 100GB hard drives will do.

I know that I'm not alone. Many people have ditched the plastic disc in exchange for the magnetic one. And anyone who has more than a couple hundred discs knows that as your collection grows, finding the track you want becomes harder. And finding new music becomes equally challenging. Why haven't computers fixed that? I'm talking about using my existing music collection to do what radio can't do.

I've given up on radio. I'm not among the naysayer patrol who believe that good music isn't being made any more. This is just a cop-out by people unwilling to put in the effort to find the stuff of quality that today's artists produce. Commercial radio does nothing to ease that process. The handful of companies that own the bulk of radio stations are just too keen on markets and demographics and a 90-minute playlist. There's not enough room for exposure and customization.

The occasional radio hit is a song that appeals to me. But most of it just isn't my cup of joe. And if I listen to the radio on my AM and PM commute, there's a damn good chance I'll hear the same Latest Hit Single twice. How is this fair? My own CD collection is a tiny fraction of what exists in the musical world, and I can go months without hearing the same track twice. But the radio, shared by thousands of people, can't even let me go an afternoon without hearing the same song three times. Can't radio be more than a 24-hour commercial for a dozen corporately approved artists?

There is a chance for the Internet to solve this. The first problem is simply software that can track your tunes. The software should scan your music collection and know what you

have available. It should monitor what you listen to. Tracking artist, album, genre, but also things like time of day, beats per minute, and so on. This data should be consolidated and sent to some sort of master server, where it is combined with that same data for thousands of other listeners.

With this data, "radio" takes on a new dimension. Instead of listening to a dozen tracks from a predetermined set list, the system can start trying to help you find songs that you want to hear from whatever audio sources are available. Mostly such a system would just play the MP3s you already have: You obviously like them. But the real magic could happen when conduits to commercial sources appear.

Put on your imagination cap and pretend that you listen to a dozen songs. Your media player compiles those dozen songs and looks up in some master database and tries to find another song that fits in with what you've heard in the last hour. Maybe this track is free. Maybe you are charged a nickel for the new MP3. Maybe it is followed by a commercial. Maybe it's \$5 a month. There are many options that I

**Can't radio be more
than a 24-hour commercial
for a dozen corporately
approved artists?**

won't get into here.

The trick is to mix and match music you own with music that might appeal to you. This keeps the needed bandwidth relatively low since you only need to fetch new tracks. The other trick is to guess intelligent songs for each listener. The odds of me buying a CD based on hearing a song on my FM tuner are almost zero. But if this new system notices my tastes and recommends something I like, the odds of me buying that CD become very high.

Traditional radio has failed us. By throwing broad demographics to the wind, listeners will be exposed to more music, and because that music will be more in tune with what they enjoy listening to, the odds of "radio" influencing CD purchases are vastly improved. It could happen. ■

Sound off to Rob at malda@cpumag.com

by Joan Wood

Hooked On Photonics



In 1980, the challenge of putting a live video camera into a race car was solved by bolting the camera lens to the car chassis (to eliminate vibration), shock mounting the body of the camera (vibration-sensitive tubes and all), and connecting them by a 3-foot-long fiber optic bundle, like a giant endoscope. I won't go into details of how the video signal got from the race car to your living room, but it involved two microwave transmitters and a helicopter circling over the racetrack.

As least senior member of the team (think water boy), one of my duties was returning the battered fiber optic bundle to the local manufacturer and exchanging it for a new one after each test run and race. I'm not sure they really understood what we were doing to their bundles, but back then, they were interested in supporting every possible use for fiber optics that walked in the door because fiber was going to change the world!

In 2000, I moved to Silicon Valley where some forward-looking genius installed fiber instead of copper between my neighborhood and the local Telco substation. Because fiber and DSL are generally incompatible, the maximum DSL connection speed I can now get is 144Kbps. This is not quite the realization of the Fiber Optic Dream the teachers raved about in school. But salvation is just down the road. (No, I'm not moving to an older neighborhood, though that would solve the problem more quickly.) While investigating Intel's plans for future CPUs (shameful plug: see the article on page 58), I got hooked on photonics.

Intel is not the only kid playing in the photonics sandbox, but with billions in research, investment, and acquisition dollars to spend, it certainly has the biggest bucket. The goal of "bringing optical networking down to the curb and ultimately to the PC" is being pursued from several directions. The prime focus is to take advantage of the cost, manufacturing, and technology advances brought about

by getting optical devices onto silicon. For example, Intel's CTO Pat Gelsinger (read our interview with him at www.smartcomputing.com/cpumag/jun02/gelsinger) demonstrated a chip at IDF (Intel Developer Forum) that functions as a software-tunable optical filter, costing about a dollar to manufacture, that replicates the functionality of a mechanical device that has to be manually reset to change wavelengths and currently sells for \$10,000.

In March, Intel launched a Photonics Business Unit to provide photonic design and manufacturing services to optical component and systems manufacturers. To service those customers, Intel recently completed a 70,000-square-foot photonic component fab facility in San Jose. Over the last few years, Intel has acquired several optoelectronics companies and invested in dozens more. So this isn't just a little side hobby. There is a reason that Intel is so interested in photonics. This excerpt from *Intel Developer Update* magazine's "Silicon Photonics" article sums it up:

Intel is not the only kid playing in the photonics sandbox, but with billions in research, investment, and acquisition dollars to spend, it certainly has the biggest bucket.

"In some ways, the photonics industry is where the semiconductor industry was in 1970. However, 30 years ago, we didn't have the Internet or the high-volume fabrication plants that we have today. We also did not have the expertise to increase microprocessor capacity, garnishing another

25MHz that we now measure in months instead of years. With photonics, circuit densities can continue to double every 18 months for decades upon decades to come." (www.intel.com/update/departments/special/ss04024.pdf)

Look for big developments in this area. Photonics-based computing, local fiber networks, more cute demos at IDF. Personally, I'm just hanging on for a decent connection. . . . ■

Flash me at joan@cpumag.com

Starting as gopher for the Emmy-winning team that pioneered live in-car TV cameras for the Indy 500, Joan became an independent video/sound engineer, technical director, and producer. Playing with Reality Engines and motion platforms led to co-founding Xatrix Entertainment where she produced the two Cyberia games. Before 3-D acceleration was trendy, she formed Mango Grits to develop hardware-only game Barrage for Activision. Since cashing out from SharkyExtreme.Com, where she was co-founder and managing editor, Joan has retired.

Road Warrior

Spring Blooms New Palms, BlackBerry Gets A Voice, A New (Legal) Use For The iPod & More From The Mobile Front

Coming To America


My biggest complaint about Sony (www.sony.com) is that it is based in Japan. As a result, the Japanese always get Sony's cool new toys first. That trend continued March 11 when Sony announced a couple of cool new PDAs for Japan that made some PDA enthusiasts in the States a little jealous. Fortunately, before I could even whine about U.S. availability, Sony announced the new PDAs were coming across the Pacific.

The Sony CLIE PEG-NR70 and PEG-NR70V have a clamshell design that opens to reveal a keyboard on one half and a 3.8-inch transfective display on the other. The display can rotate 180 degrees, letting you close the device with the display facing out.

The CLIEs include more than just a new design and built-in keyboard. The transfective display is a combination of older backlit TFT displays and the newer reflective TFT displays that provide better outdoor viewing. A virtual Graffiti area occupies the bottom third of the display, but certain applications can eliminate it. Sony's Picture Stand application, for instance, can hide the Graffiti area to show pictures, using the entire 480 x 320 display. Sony will provide information to developers who want to use the virtual Graffiti area in their applications.

Like other high-end CLIEs, the new models support digital audio but also include a polyphonic speaker that lets you listen to music without headphones. The new CLIEs are also the first Palm OS devices to run the new 66MHz Dragonball Super VZ processor. It shouldn't surprise anyone that both models include a Memory Stick slot. ▲

Sony's new CLIEs feature a clamshell design and a built-in keyboard. The PEG-NR70V shown here includes a digital camera that lets you take snapshots you can save in your CLIE.



Apple's New PDA: The iPod

Every year around the time of Macworld, rumors start flying about an Apple PDA. Pictures and even video have turned up that reportedly depict some sort of Apple-branded PDA in action. Indubitably, the Jobs' Macworld keynote comes and goes and the latest roundup of "insanely" great products once again does not include an Apple PDA. The rumormongers then start working on next year's gossip.

As it turns out, Apple already has a PDA on the market; it's called the iPod. Users are finding clever hacks for the versatile little device and its 5GB hard drive, from the criminal to ingenious. ProVUE Development (www.provue.com) is on the ingenious side of the fence with its Panorama iPod Organizer, a nifty little program that lets you store personal data on your iPod. The software lets you access names, phone numbers, e-mail addresses,

appointments, flight numbers, and a host of other data. The iPod wasn't designed for data entry, but it's still nice to be able to retrieve data when you need it. The iPod Organizer software is available for \$19.95. ▲

You can make Apple's iPod a little more versatile with ProVUE Development's iPod Organizer software. Using the software, you can import personal data from your Mac to your iPod.



RIM Speaks Up

RIM (www.rim.com) announced a new product in early March, the BlackBerry 5810. It looks much like the BlackBerry 957 but includes one major new addition: voice support. In addition to receiving e-mail, you can also receive phone calls and SMS messages.

RIM didn't change much in terms of the 5810's design. The unit doesn't look anything like a phone and uses an ear bud for voice communications. Although this looks awkward, it has the benefit of letting the user read the display and use the device while talking on the phone. For example, if you need to check your calendar while you're on the phone, the 5810 will let you.

In the United States, VoiceStream and AT&T will provide service. ▲

Can You (O)play?

MPEG-4 promises to deliver streaming video to low bit-rate devices, such as wireless clients. But you probably already knew that. What you may be less familiar with is an upstart challenger to MPEG-4. Developed by Oplayo (www.oplayo.com), the new technology uses MVQ (Motion Vector Quantization) to deliver tightly compressed video to small devices over slower connections.

According to Oplayo, the company can already stream video over a 30Kbps connection. A related product, the oplayo compressor, uses MVQ technology to produce small video clips. You can view content produced with the oplayo compressor on a desktop without additional software players or plug-ins. ▲

The m130 and m515 are the latest PDAs from Palm. The m515 will replace the m505, and the m130 is the lowest-priced color device on the market.



Spring Brings New Palms

Who doesn't love springtime? The green grass, warming temperatures, and of course, new models from Palm (www.palm.com). This year, Palm's spring lineup includes the i705 (see May's Road Warrior for more information) and the m130 and m515. In sticking with Palm's recent and infinitely more rational naming practices, the m130 represents a low-end, entry-level device while the m515 is the latest premium offering.

Most of you probably remember the flack Palm got last year when it introduced the m505. Complaints about the brightness of that model's display flooded message boards on Palm-related sites. Palm has reportedly corrected this problem with the m515. Not only is the sidelight brighter, but users can also now actually adjust the brightness level (as opposed to using the On/Off brightness settings on the M505). The m515 also includes 16MB of RAM to accommodate the needs of high-end users, such as those wanting access to corporate databases with their Palms. The biggest news is probably that the m515 will completely replace the m505. The m515's price will be the same as the m505's at \$399.


Despite the improvements in the m515's display, the m130 might make the biggest splash this spring. The entry-level system includes a rechargeable battery and color display for \$279, making it one of the lower priced color devices on the market.

Like Sony's new CLIEs, the m130 uses a transreflective TFT display that uses a backlight to illuminate the display rather than a sidelight. The display can still reflect ambient light, making it easy to see outdoors. The m130's display is passive matrix, meaning it will update slower than the active matrix display on the m515 and other PDAs.

In other Palm news, the company released its Bluetooth Card (\$129) in March. The card is compatible with any Palm device with a SD card slot and lets your Palm communicate with other Bluetooth-enabled devices, including mobile phones, notebooks, and other PDAs. BlueChat and BlueBoard applications are included with the card and let users communicate with one another. Using BlueChat, for instance, you can select another Palm user and secretly beam notes back and forth. BlueBoard is similar but is more of a whiteboard environment where a user can draw directly on her PDA display and share information with another user. ▲



At Your Leisure



Plug In, Sit Back & Fire Away

The entertainment world, at least where it pertains to technology, morphs, twists, turns, and fires so fast it's hard to keep up. But that's exactly why we love it. For the lowdown on the latest in game consoles, games, PCs, DVDs, and just stuff we love, read on.

Star Wars Extravaganza

In honor of the impending theatrical release of "Star Wars Episode II: Attack Of The Clones," we've put together a small selection of recent "Star Wars"-themed games from LucasArts.

Jedi Knight II: Jedi Outcast Star Wars Action At Its Best

Jedi Outcast (developed by Raven Software) has Kyle Katarn returning after stints in the previous two well-regarded Star Wars FPS games. You'll spend the first several levels fighting through waves of storm troopers. The good news: Storm troopers will try to flank you, they'll retreat under heavy fire, and their death animations are

Once you get your lightsaber and begin using the Force, Jedi Outcast becomes a whole new game. Your developing skills with the Force and the lightsaber improve through the levels, and you begin to feel and look just like a Jedi from the movies.

After you're done with the single-player game, there's a lot of fun to be had with the numerous multiplayer modes. You can even go one-on-one with a friend using



The first five levels will have you mowing through swarms of storm troopers. Bear with it.

just your lightsaber. Jedi Outcast comes highly recommended to all action fans.

Star Wars Jedi Knight II: Jedi Outcast (PC)

\$49.95 (\$59.99 for the Collector's Edition)

LucasArts

www.lucasarts.com/products/outcast



You have the power to wield a lightsaber like a true Jedi Knight. Two words: Bloody amazing.

"Star Wars" through and through. The bad news: It just gets monotonous after awhile. Fortunately, the first five levels don't represent the entire game; the remaining 19 levels of the single-player mode are easily worth the full price of admission.

Check These Out On The Web

See our reviews of **Sega Soccer Slam** (NGC), **Tiger Woods PGA Tour 2002** (PS2), and **Home Run KING** (NGC) at www.smartcomputing.com/cpumag/jun02/gamereviews.



With its colorful characters and a variety of signature moves, Sega Soccer Slam is the NBA Jam of soccer games.



Even playing as Tiger, you may have to dig one out of the sand once in a while in EA's Tiger Woods PGA Tour 2002.

Tino Martinez stretches to snag a fly in



Sega Sports' Home Run KING.

Star Wars Rogue Leader: Rogue Squadron II The Force Is Strong With This One



If you've played the original *Rogue Squadron* on the N64, you'll be instantly familiar with the sequel's play mechanics and objective-driven missions. The controls are similar, as well, although the GameCube's quirky control pad forces you to learn some new tricks while piloting

Rogue Leader puts you in the cockpit of a Rebel Alliance B-Wing for an assault against an enormous, wonderfully detailed Star Destroyer.

a number of Rebel Alliance starfighters. The largest distinction between this game and its predecessor is the jaw-dropping quality of its graphics. The original looked good, but the GameCube brings the "Star Wars" universe to life in a special way, thanks to crisp textures and nice touches such as spot-on light sourcing.

Also of note is *Rogue Leader's* slightly higher learning curve. LucasArts listened to gamers' cries for space battles, but fans of the original will quickly find that the lack of firm visual reference points makes the game a little tougher at the outset.

We're glad LucasArts included space battles, though, and they're well worth the extra playtime it takes to learn the ropes. If you've seen many previews/reviews of *Rogue Leader*, you probably know about the epic battle to save a Rebel frigate from an Imperial Star Destroyer that takes place just a few missions into the game. This mission alone is probably worth the price of the game for rabid Star Wars fans, and there's a lot more game left when you finally survive it. The addition of multiple wingmen that follow your commands and a targeting computer both mitigate the game's difficulty and add a lot of fun.

Rogue Leader isn't a revolutionary game, but it's probably the best GameCube-exclusive title on the shelves this spring.

Star Wars Rogue Leader: Rogue Squadron II (NGC)

\$49.99

LucasArts

www.lucasarts.com/products/rogueleader

Star Wars Jedi Starfighter There's A New Jedi In Town

SWJS picks up where *Star Wars Starfighter* left off, dropping you into the role of two characters that fly numerous starfighter missions as the events of Episode II unfold. The game's primary protagonist, Adi Gallia, is a Jedi Master instructed by Mace Windu to put a prototype Jedi starfighter through extensive combat testing. In addition to the ship's standard complement of weapons and defenses, you'll be able to up the ante a bit with help from Adi's Force powers.

Said Force powers include four main techniques (Force Shield, Force Lightning, Force Reflex, and Force Shock Wave), each with three modes, or effective states. You'll begin with access to just the Force Shield and will unlock the others as the game progresses.

SWJS' other main character, Nym, is a not-so-cute alien who commands a small band of pirates in a perpetual life-and-death struggle with the Trade Federation. Nym flies a Havoc, which has heavier shields and

more weapons to choose from than the Jedi starfighter but of course lacks Force powers.

Like its predecessor, *SWJS* puts you in the cockpits of its characters' fighters for a variety of missions, ranging from simple training exercises to pitched battles with



The Delta-7 is a small but agile craft designed specially for the Jedi, protectors of order throughout the Republic.



Jango Fett's Slave I unleashes a barrage of hot laser death.

Trade Federation land, air, and sea forces. The game's controls are relatively easy to master, making it accessible from the

moment you pick up a controller. As you progress, missions get harder, and you obtain the use of more force powers/weapons, but the real hook is *SWJS'* intriguing storyline. You'll come across Yoda, Count Dooku, Jango Fett, and a number of other Episode II personalities as you team up with Nym's pirates to foil a plot by Dooku's separatist group and the Trade Federation to take control of the Karthakk system.

OK, so maybe that's not as engrossing as leading a squadron of X-Wings on an attack against the Death Star, but the game is a lot of fun to play and will do wonders for that Star Wars jones you've got going.

Star Wars Jedi Starfighter (PS2)

\$49.99

LucasArts

www.lucasarts.com/products/jedi_starfighter



The DVD Bit: Only You Can Prevent Full Frame

For home video enthusiasts, it all began with laserdisc. The LP-sized video format catered to film connoisseurs, giving a full-on view inside the filmmaking process and historical look behind our favorite films. And everything that people loved about laserdisc, they found on the pint-sized, but surprisingly bigger DVD format. Commentaries, deleted scenes, digital surround sound . . . they were all there, with new features such as anamorphic (www.thedigitalbits.com/articles/anamorphic) widescreen video to sweeten the deal. The greater storage capability and digital nature of DVD allowed for even more possibilities such as more language tracks, parental blocking, and alternate angle features. And the smaller disc size made it easier to package multi-disc sets with even more special edition material. All of these things attracted consumers. And wherever consumers go, marketing departments follow.

In this new, more lucrative world of home video, special edition material has suddenly become "added value" material. Blockbuster, which initially shunned DVD, now sees a rental goldmine. But as the format grows in popularity, more people have been complaining: "Why are there black bars on my television?" As odd as it is to hear, widescreen DVDs are being returned as defective, which gives reason for Blockbuster to push the studios for more

full-frame versions. When these become available, stores stock fewer widescreen titles. What none of these people realized is,

The best thing you can do is educate your friends and family about the benefits of the DVD format and the benefits of anamorphic widescreen enhancement.

by cutting up their picture to a standard TV-friendly, full-frame size, they're missing a large portion of the movie. The nuances, the conversation style, the composition of the frame—most of what makes a great filmmaker great, are being chopped away so viewers don't have to get used to the "black bars." Instead of educating people on what the black bars really represent, they take a world carved out by home video fans for over a decade and start filling it back in with ignorance.

And those "added value" materials that film fans love are slowly but surely becoming fluffy and less interesting. Why? Well, as Hollywood stars watched DVD grow, they saw the studios begin raking in money, and they demanded their share of it. The last Screen Actors Guild strike had more to do with shared DVD revenue than anything else. And as the actors' share of the DVD profit pie increases, the studios struggle to keep their own slice unchanged. That usually means cutting corners; "added value" material adds less value all the time. It's a win-win for Hollywood, but all those film connoisseurs that made DVD a household name in the first place are left hanging.

But all is not lost. DVD still provides connoisseurs, and the rest of us, with the highest available video and audio quality for watching films at home. Meanwhile, new formats such as D-VHS



A side-by-side look at Paramount's *Forrest Gump* DVD next to its laserdisc predecessor.

DVD Byte by Todd Doogan

Fox has a great idea here. Take an incredibly well-made thriller just



made for DVD and release a two-disc set. I know that doesn't sound too incredible, but get this next part: The two-disc set will be available only for a limited time. Once it's gone, the super special edition of "From Hell" will get pulled, and a single-disc release will come out later this year. Disc one is packed with an in-depth commentary



track and a large selection of deleted scenes. The limited second disc has a featurette on Absinthe, a tour of the Ripper murder sites, and a very thorough documentary made exclusively for this DVD with archival footage littered in via a "white rabbit" feature. ▲

and HD-DVD loom far away on the horizon. Should we all wait for these formats and stop buying DVD? No. DVD is as good as we're going to get for a while. These newer technologies, although inherently better, are too far away logistically, technically, and financially. The best thing you can do is to educate your friends and family about the benefits of the DVD format and the benefits of anamorphic widescreen enhancement. That way, when HD-DVD does come around, we might all know what those black bars are, why they *should* be there, and how it's all a good thing.

by Todd Doogan

Hot Shots: The Beauty Of The Game

Yeah, we know it's all about the gameplay. Who needs eye candy, right? (Or so grognards and some "old school" gamers like to moan.) Sure, we loved M.U.L.E., Elite, Ultima, Balance Of Power, Civilization, and Dune as much as the next guy. But we also know there's nothing better than stellar gameplay combined with knockout graphics to make that killer game. These games all have potential.



No One Lives Forever 2 (PC). Say goodbye to 007. Secret agent Cate Archer will be back looking better than ever using the new LithTech 3D engine. The scenes you see behind Cate are CPU mag exclusives that show off some of the new environments. Monolith Productions (www.lith.com) is hard at work getting this anticipated sequel into stores by the end of the year.



Infinite Loop

Serious Dough For A Sweet Ride

Scooters are supposed to be a cheap mode of transportation, right? Not when the scooter in question is a Segway HT (Human Transporter). Commercial versions of the Segway HT are supposed to appear later in 2002 or early in 2003 with an expected retail price of about \$3,000. However, inventor Dean Kamen decided to auction off three "preview" consumer models to raise money for his F.I.R.S.T. charity. The top bid for a Segway HT was \$160,100. The other two models went for \$104,100, and \$100,600. That's a lot of script for a sidewalk jockey to pony up for a sweet ride. Compare the price of one Segway HT to these more affordable modes of transportation:

Mercedes SL500: \$104,320


Corvette Z06: \$51,445

Dodge Durango SLT Plus 4X4: \$37,905

Harley Davidson Softail Deuce: \$16,555

Honda Metropolitan scooter: \$1,699

Nike Air Max Healthwalker IV sneakers: \$75



The Y Project (PC). Step forward 200 years in this sci-fi FPS title that's awash with adventure and RPG elements. German developer Westa Interactive has enhanced the latest version of the Unreal technology engine (not publically available at press time) with plans to have the game poised for release by Q1, 2003. Learn more at www.y-pro.net.

\$\$\$\$\$

CONTROL ACCESS

Whip Your Work Into Shape

LIKE ANY OTHER OFFICE PROGRAM, YOU CAN CREATE MACROS IN ACCESS THAT SPEED UP REPETITIVE OPERATIONS (SUCH AS CREATING AND PRINTING MONTHLY SALES REPORTS).


Even better, you can attach macros to buttons to make it easier for users to get around forms and give you, the developer, more control over how the form is used.

To create a macro, click the Macros button in the Objects section of the Database window. Choose the New button at the top of the window to open the macro in Design view. Now assign some instructions (actions) to the macro.


Click in the first empty row in the Action column of the macro and then click the drop-down list arrow to display a list of preset actions. Click an action to add it to the list, keeping in mind that Access will go through the actions for the macro in the order you set. Depending on the action that you choose, Access may require action arguments so that it knows exactly how to carry out your instructions.

Let's create a macro to open a table and then tell the user to get to work. Just for interest (and to get people's attention) we'll add a beep to the macro, too. Choose OpenTable from the drop-down list. Because this action has arguments associated with it, related arguments automatically display at the bottom of the Macro window. Click in the first argument for the OpenTable action (Table Name) and then glance at the information in the lower right corner of the window. This information summarizes whether the argument is required

and gives you a few hints on how to set it up. Click the drop-down list arrow in the Table Name box and then choose which table you want to open. Alternately, you can drag the table object from the Database window into the Macro Design window.



Don't sweat creating controls on forms. Let the Command Button Wizard do the work for you.



If you're tired of looking at the same old thing, you can add a picture to your form.

Now you can add additional arguments, all of which are designed to give you control over how the database is used and what will happen when the macro runs. For example, you can

choose the view (such as Datasheet, Design, or Print Preview) you'll let others use by choosing it from the View drop-down list. Additional arguments help you exert your power over the way the database is used, such as setting the Data Mode to Read-Only.

Add more commands. When you have put the finishing touches on the first action, add some more commands, keeping them in the order that you want them to execute. Specify Beep as the next action (an action that does not have any arguments to set, incidentally). Then choose MsgBox as the last action so that you can splash a message on the screen. MsgBox does have arguments, including the comment you want to appear in the box. To add the text you want, click in the Message box in the Action Arguments pane and then type **Get to work! The table is open!** Additionally, you can set the Type of message box (such as Critical or Warning) and even create title bar text (I'm watching your every move) by entering it in the Title box.

To test the macro, save it and then click the Run button. If all goes well, the table will open, the message will appear on-screen, and the beeping sound will rouse the co-worker dozing in the desk next to you. If you need to edit the macro later, select it in the Database window and then choose Design. Things you may want to do to tweak the macro? Add more actions (using the Insert Row command if necessary to create space between the existing actions), rearrange the order of actions by dragging the Row Selector(s) to a new location, or change the arguments for specific actions. Of course, you can also select and then delete actions on the macro list, as well.

Command Buttons

After you've developed a few macros, you're ready to take control of your forms by adding command buttons (with macros) to them. Command buttons, just as their name implies, give you a painless way to perform actions simply by clicking them on the form. When you click the button, it carries out the associated command sequence, such as opening a table, printing a report, or sorting records.

There are two ways in which you can create a command button: from the ground up or by relying on Access' Command Button Wizard to lay the foundation and then revising the associated properties to fit your needs. Using the wizard speeds things up because you can use one of the more than 30 built-in command buttons that perform actions such as finding a record, running a query, or printing a report, but developing the button from scratch gives you more flexibility and is a challenge most of you will enjoy.

The wizard. Take a look first at using the Command Button Wizard. Display your form in Design view. If the Toolbox (with its set of controls) isn't displayed, choose View, Toolbox. Use Screen-Tips to identify and then select the Control Wizards button. This gets the wizards geared up so that any control you choose in the Toolbox will automatically activate the associated wizard.

Now launch the Command Button Wizard. Click the Command Button icon on the Toolbox and then click in your form where you want the button to appear. Access will automatically create a default-sized button and then start the Command Button Wizard.


Right away you'll notice that you have a number of main categories from which to choose; click one of them to see the actions that are available for each. Just for fun, choose the Miscellaneous category and then click

Print Table before advancing to the next wizard screen. Choose one of your database tables on the list. When you advance to the next screen, you'll have a choice of what you want to appear on the button: text or picture. Finish off

Hide Your Handiwork

After you've laboriously created your macros, you probably don't want other people messing with them. So hide the macros from view without deleting them; if others do wander across them in the Database window, they can't delete or change them.

To hide a macro from view, right-click it in the Database window and then choose Properties. In the Macro Properties dialog box, check the box for Hidden before clicking OK. To view the macros again, choose Tools, Options and then click the View tab. Check the Hidden Objects box and then choose Apply.



See the macro? Neither do we, but that's the point. If you hide your macros, it'll be harder for others to mess with them.

the process, and a command button is created on your form. Like other objects, you can move and resize the button so it appears as you want. And remember, the button becomes active in Form (not Design) view.

If you want to see the VBA code created through this process, right-click

the button in Design view and then choose Properties. In the Command Button dialog box, make sure the Event page is shown. Right-click [Event Procedure] in the On Click box and then choose Build to view the code associated with your button.

From scratch. Of course, while using the Command Button Wizard is speedy, you can exert much more control over the entire process by creating the button from scratch and then attaching a macro to it. Here's how: Assuming you've already created the macro, display your form in Design view and turn off the Control Wizards button on the Toolbox. Click Command Button and then click in the form where you want the button to appear.


Right-click the button and choose Properties. In the Properties dialog box, display the Event page. This page controls the actions that are associated with the button. If you want the action to occur when the mouse is clicked, for instance, click in the On Click box and then choose a macro from the drop-down list.

If you forgot to record the macro ahead of time, all is not lost: you can create it on the fly. In the On Click box, click the Build button (the one with the three trailing dots) and then double-click Macro Builder in the Choose builder dialog box. Save and record the macro that you want to attach to the button. Finally, you can add a caption or picture to the button using options in the Format tab of the Properties dialog box.

Combo & List Boxes

Another great way of controlling your users (OK, OK, and helping them out a bit, too) is to create a combo box on a form. For example, you may want to allow your users to only enter three states in the State field: Virginia, New York, and California. Here's how:

To create a combo box, display your form in Design view with the Toolbox




It won't take long to frustrate your friends when you start attaching messages such as this one to macros—especially when you add a beeping noise that sounds every time they run the macro.

active. Also activate the Control Wizards button in the Toolbox. Click the Combo Box button and then click your form to activate the Combo Box Wizard. Choose the I Will Type In The Values That I Want radio button before advancing to the next wizard screen. On the next page, type in the first value (Virginia) and then press TAB (not ENTER) to create another row. When you are finished entering values, advance to the next screen and store the value in the State field. Optionally, on the next screen you can create an easy-to-understand label for the combo box, such as “Choose a state from the list” and then finish off the wizard. To test how well your combo box works, display the form in Form view and choose values from the combo box's drop-down list.

If you want to help users enter specific values in a field but have fairly limited choices, you can use a list box instead of a combo box. The list box displays the choices in list form, and it adds a vertical scroll bar if the list gets too long.

You'll follow the same basic procedure to create a list box as you did to create a

to use a form for promotional purposes, it's important to make it look good. And because you have the tools to do it, why not?



Make it easy to run macros by attaching them to command buttons on your forms.

Display the form in Design view and then click Image on the Toolbox. Click the form to create a default-sized image; click and drag if you want to specify an exact size. In the Insert Picture dialog box that automatically displays, choose a photograph or other picture and then choose OK. Like other objects, you can move and resize the image so that it appears where you want it. To make it appear as the background for the entire form, select it and then choose Format, Send To Back.

For even more control over how the image appears on the form, right-click it and then choose Properties. In the Format tab, adjust the properties as you wish. For example, you can choose Screen Only in the Display When box to have the picture appear when people are editing on-screen but not when the form is printed. **CPU**

by Linda Bird

Infinite Loop

What's 364.4 Smoots & An Ear Long?

MIT has a long-standing tradition of ingenious pranks. One such prank has become an integral part of local lore. In 1958, the Lambda Chi Alpha fraternity pledgemaster devised a plan to mark off portions of the Harvard Bridge connecting Boston to Cambridge using a pledge. He decided to use the shortest pledge that year, Oliver R. Smoot. Smoot's fraternity brothers used him as a human ruler, placing marks on the bridge every 10 Smoots. They discovered the bridge was roughly 364.4 Smoots and one ear. (The ear was thrown in to account for any small errors.) The marks are repainted every year and have become an accepted marker; even local police use them in accident reports.

We thought that it might be fun to calculate some measurements using Smoots.

- Smallest distance between the orbits of Earth and Mars: 34,046,208,000 Smoots.
- If every person on Earth were the same size as Oliver R. Smoot Jr., we wouldn't even have enough people to reach one-fifth of the way to Mars.
- San Francisco Golden Gate Bridge: 1,608.634 Smoots (approximately four times the size of the Harvard Bridge).
- Length of a football field (goal line to goal line): 53.735 Smoots (a little more than one-tenth the size of the Harvard Bridge).



tomorrow
morning,
will you choose
to give 110%?
75%?
33%?
14%?



Meet the newest Harvest® bar. The delicious morning energy bar from PowerBar. Made with sweet raisins, whole oats, and dipped in yogurt. Plus it's loaded with soy protein, and 16 essential vitamins and minerals for the energy to help you power through your morning. So grab one. And give it your all. Maybe even a little more.



Be great.

KILLER HARDWARE TIPS

Too Cool To Follow Rules

IF YOU'RE A USER WHO CAN'T LEAVE THINGS ALONE IF THERE'S A CHANCE YOU CAN MAKE THEM BETTER, EVEN IF IT MEANS BUCKING THE NORM, READ ON. WE HAVE

some tips and tricks that just might make you more productive.

Old Hardware Lives On

Computer magazines (*CPU*, too) devote countless pages to extol the latest and greatest hardware. Perusing these pages without feeling your month-old 2GHz screamer is already obsolete is hard. Whether your hardware is a week or decade old, take heart. With some care, you can extend your PC's lifespan for a long, long time.

It doesn't take a rocket-like PC to surf the Web, do word processing, and perform other pedestrian computing tasks. Even an old Pentium 75 can do those jobs nicely—with a bit of patience and preparation.

The right software is the most important asset for extending the life of older hardware. When you're using a slow machine with limited RAM and drive space, there's no room for bloated software. Forget about WinXP and Microsoft Office. Think thin.

Miles Wolbe, a computer consultant in Kula, Maui, Hawaii, runs TinyApps.org (www.tinyapps.org), a Web site that distributes sleek, efficient Windows apps. The site defines "tiny" as any app that fits on a floppy and doesn't require installation—a breed of programs many users have written off as extinct. The site lists more than 100 programs averaging 226KB in size, most of which are free. You won't be stuck with abandonware

from 1992, either. Almost all of the applications at TinyApps.org are regularly updated.

Wolbe's favorites include QNX Demo (www.qnx.com/iat)—which includes a Web browser, Web server, POSIX-certified real-time OS, TCP/IP stack, and a windowing system—and Off By One Web Browser (www.offbyone.com), a small, self-contained Web browser with full HTML 3.2 support.

"Both of these represent quite a paradigm shift for users accustomed to 1GB-plus operating systems and 25MB-plus



The Sound Blaster Audigy: audio gadget of the future or rehashed technology? The truth is out there—somewhere.

browsers," Wolbe says. "Instead of continuing to throw hardware at the problem, these programmers have decided to do something radical: write good, clean code."

These sleek applications are instrumental in stretching the lifespan of PC hardware, Wolbe says. "I just salvaged a 166MHz desktop computer, loaded it with Windows 95 and some tiny apps. It boots in about 20 seconds and handles Internet and word-processing tasks beautifully. Perhaps even more incredibly, the

whole system was completely free at a computer recycling day."

Putting to use the old (but still functional) dirt-cheap computers you find at garage sales, thrift stores, and in classified ads may be the biggest benefit of tiny apps. If you salvage an old PC and monitor from a trash heap, many Web sites (try www.computergeeks.com and www.computersurplusoutlet.com) offer inexpensive accessories, such as 30-pin SIMMs, network cards, and cheap hard drives, to spruce them up.

You don't even have to pay for an OS. Visit Linuxiso.org (www.linuxiso.org) for a selection of Linux installation ISO images ready to burn to a CD-R. You can also download FreeDOS, which includes a graphical windowing system, at www.freedos.org.

You could also turn that PC into a router and firewall for your LAN by installing a couple of Ethernet cards and downloading the Linux Router Project (www.linuxrouter.org), which is a micro distribution of Linux that is dedicated to networking.

Also, be sure to visit Low End PC (www.lowendpc.com), a Web site offering articles and tips on using older PCs. A sister site, Low End Mac (www.lowendmac.com), is dedicated to keeping Macs out of the graveyard.

The Truth Is Out There

An alien autopsy took place at Area 51. Elvis is alive. The Mafia killed JFK. The moon landings were staged. Here's a new one to add to your list of conspiracy theories: Creative Labs' Sound Blaster Audigy sound card is just a repackaged Sound Blaster Live! with new drivers.

At least that's what a group of Taiwanese hackers were alleging. The group claimed that by modifying the drivers, any Sound Blaster Live! card can take on the full functionality of the Audigy,



with the exception of the Audigy's IEEE 1394 (FireWire) port. The Sound Blaster Audigy offers 24-bit, 96KHz sound, bringing PC acoustics into the realm of high-end consumer electronics audio.

In the spirit of all good conspiracy theories, one vital piece of evidence is difficult to obtain. Performing the hack requires the SB Live! 2 SB Audigy driver, but the sites that allegedly house that crucial file are flaky at best and gone forever at worst. Users who have managed to acquire the drivers claim the hack does work, although some have reported a few technical glitches.

"I posted the story back in August 2001 [since taken down], and at that time the first people to bring this to the attention of the public was AudigyBoom. They leaked the original SB Live! 2 SB Audigy Driver," says Bob Kahlon, an editor at Warp2Search.net (www.warp2search.net) who goes by the handle "News-Factory." AudigyBoom is known as the kX Project (www.kxdev.com).

Does Kahlon think the Audigy is really the Sound Blaster Live! repackaged? "I would say it's possible. The only people that really know are Creative and maybe the kX Project," says Kahlon. There is precedent. Last year, digit-life.com (www.digit-life.com) released software that would make a four-channel Sound Blaster Live! card support 5.1-channel audio.

Also in the spirit of good conspiracies, The Man completely denies the allegations.

"That is not even close to true in any respect," says Phil O'Shaughnessy, public relations manager for Creative. "The Audigy is an entirely redesigned audio architecture. It has four times the processing power of the EMU10K1 products," the processor in the Sound Blaster Live! "The Audigy processor has EAX advanced functions, such as environmental morphing and environmental panning, that are not possible with the Sound Blaster Live!

The Audigy achieves 24-bit, 96KHz,

100dB signal-to-noise ratio that cannot be achieved with the Sound Blaster Live!"

We don't know if an alien cadaver is hidden somewhere in the Nevada desert or if the Audigy hack is for real. We do know that the truth is out there—somewhere.

Toys You Can Interface

Toys are fun. Toys that your computer can control are even better.


Lego Mindstorms (mindstorms.lego.com) is perhaps the quintessential computer-connectable toy. You can build robots using old-fashioned Lego building bricks, then program and control them with your PC. Data is sent wirelessly from an infrared transmitter to a special programmable Lego block. A massive online community has been built around Mindstorms, but it's only the tip of the PC toy iceberg.

The Boe Bot Robot Kit (www.parallaxinc.com) is a more "serious" toy that should satisfy students learning about robotics. You can build a rolling robot (\$185), which includes infrared object detection and

photo-resisters, in about two hours. In addition, you can program the Boe Bot to follow a line, solve a maze, pursue a light source, or communicate with another robot. More advanced roboteers will appreciate the \$219 GrowBot, which takes about five hours to assemble and is more expandable.

Toys that talk to the PC aren't limited to robots. LeapFrog (www.leapfrog.com) offers a line of kid toys that can download data via a computer. Turbo Twist Spelling teaches spelling skills to kids in grades 1 to 8. You can download grade-appropriate spelling words to the toy from LeapFrog's Web site or create and download your own spelling lists.

Aimed at middle-school students, LeapFrog's iQuest handheld is a kid-friendly PDA with scheduler, address book, dictionary, and calculator functions.



Download this week's math chapter into LeapFrog's iQuest handheld. Just don't use it during the test, kids.

LeapFrog has partnered with textbook publishers to offer downloadable curriculum, including students importing chapter outlines and interactive quizzes for their math, science, and social studies classes.

If programmable toys don't interest you, consider hacking a toy not meant for programming. Hacking Furbys (www.afu.com/fur.html) to speak fluent French or do the Macarena is yesterday's news, but how about modding your Sony AIBO? AiboPet's Reborn Site, (aibohack.com) offers programs that will change the personality of the AIBO. You can make it disco, behave like Cartman, speak German, or even act like a kitty.

Other hackers have reprogrammed the Cybiko (check out www.devrs.com/cybiko), another handheld computer for kids, and those obnoxious Boogie Bass and Big Mouth Billy Bass singing fish (www.ai.mit.edu/~vona/bass/bass.html). **CPU**

by Kevin Savetz

Infinite Loop Tons 'o Data

Here's why data storage is such a growth business: Humans produce lots of information. (Just think about that closet you have that's full of old LPs, tax records, 8mm home movies, and the like.) Computers and the Web have just accelerated the trend. According to research from the University of California, Berkeley, (www.sims.berkeley.edu/research/projects/how-much-info/internet.html) the world now produces about 2 exabytes (1 exabyte = 1 billion gigabytes) of information each year. How much is that? Well, all the words ever spoken by human beings would be only the equivalent of about 1 exabyte. We need more closet space...



Technically Speaking

An Interview With Greg Bildson, CTO & COO Of Lime Wire



When the Recording Industry Association of America sued Napster, it set off a legal firestorm that has left most peer-to-peer services similarly besieged. Morpheus, KaZaA, Scour, and others have either faced lawsuits or gone out of business. However, LimeWire, a P2P client Lime Wire (www.limewire.com) developed, remains removed from the RIAA's wrath—at least for now.

The argument for LimeWire goes that P2P apps don't pirate software, people pirate software. Illicit copying of protected media is an abuse of the platform that designers never had in mind. What makes Lime Wire's case interesting is that the company actually has a history and product to back up this claim.

We spoke with Greg Bildson, Lime Wire's CTO and COO, about the company's technology and its position in the digital music war and how its platform may ultimately transform the way we share information.

by William Van Winkle

CPU: Let's start with a look at LimeWire's roots. Where does the product and company come from?

Bildson: In early 2000, the Lime Group was investigating new technologies and ideas related to the Internet. I was working on a project called Lime Objects, which was focused on distributing rich product information between businesses. While that was going on, we were watching the file sharing space, things like Freenet and Gnutella. We saw those as interesting ways to distribute information to companies and just interesting technology in general. Our owner saw that this form of distribution, querying and responding with information across multiple sources of that information, was the way of the future. That's why we got into the Gnutella space. It really wasn't because of the file sharing. Certainly it wasn't because of anything music-related. We basically saw this as a new way to distribute information and access information.

Our goal is to make Gnutella a very popular Internet protocol, and on top of that, allow the general public to tie into not only other members' information, but also into various corporations. To do all of that, though, we need to make Gnutella a very fundamental, ubiquitous Internet technology. That's how we realized that to move Gnutella in this direction, we needed to create a Gnutella client. We did that in November of 2000, and it became very popular. We hit 300,000 daily users in 2001. Over time, though, the client took up a lot of our efforts. When it came down to crunch time and we actually needed to generate revenue, we realized that our long-term business—providing servers and services to corporations—was not ready to generate revenue. We realized that the client is where we could actually

generate revenue in the short-term by selling a LimeWire Pro version and doing advertising on the client.

We're still working on this long-term business model of adding transactionality and corporate information into the network, but that's going to take a fair amount of time to evolve.

CPU: What is the fundamental difference between the Gnutella network LimeWire runs on and Napster's old technology?

Bildson: Gnutella is a totally distributed network in that there are no central servers. Napster had a central server that indexed all of the users on the network, and your searches only went to the Napster server. Only the uploads and downloads were done in a distributed fashion. On Gnutella, everything is done in a distributed fashion. The central server's job has been distributed to every individual user on the network, although there is now some degree of indexing capability within a local neighborhood.

CPU: What is the difference between the free LimeWire and LimeWire Pro versions?

Bildson: Basically we have an advertising banner at the bottom and we take that out on the Pro version. The Pro version also has e-mail support, and we don't bundle any additional software, like third-party software packages or Desktop icons.

CPU: LimeWire no longer uses, to use the common term, "spyware?"

Bildson: Right, like Gator. We've made an effort to take out those apps.

CPU: You're making an effort, but are they in or out?

Bildson: They're gone from the free version.

CPU: Will they come back?

Bildson: Hopefully not. It's contingent on us making money on the Pro version, our little icon drops, and things like that we do instead. That kind of advertising is what basically pays for the free version.

CPU: How much revenue does LimeWire generate?

Bildson: I don't think we can report on numbers, but let's say the company is operating in the black.

CPU: It seems probable that LimeWire is most often used for transferring audio and video files, but in your mind, what should people be using LimeWire for?

Bildson: [long pause] I mean, they should use it based on the terms of use, which tell them not to share restricted or copyrighted materials. Basically, all I can say is that enforcement should be done at the user level, and it should also be done in corporations. No one is going to go out to Microsoft and say, "make sure that your e-mail products don't distribute certain types of files, make sure that your Web servers don't distribute certain types of content, and make sure that your FTP servers don't distribute certain types of files." I don't think you can just outlaw specific types of technology.

Our desire is to build capabilities for stuff like a global knowledge network where you can actually search for any kind of information that you want, where the combined intelligence of the world is available online to you in a fashion that is richer than the Web. So if a biologist is searching for the latest cancer research on a particular part of the genome, he can go find that. I think there are a lot of extremely beneficial uses that this technology can provide in the world. It would be extremely shortsighted to cut it off for the sake of making sure that certain limited interests are satisfied.

CPU: Is it a fair statement to say that Napster was shut down because the company would not or could not restrict the flow of copyrighted material across its network?

Bildson: They were definitely asked to restrict the content on their network.

CPU: But they couldn't, so the RIAA took Napster, and a lot of others, to court. But I haven't heard of any suits against LimeWire. Is the company under any suit right now?

Bildson: No.

CPU: How did you dodge that bullet?

Bildson: No comment.

CPU: Are the RIAA's actions fair? What is the solution?

Bildson: I don't know what the solution is, but we do believe that copyright enforcement is a legitimate thing. But it should be done at the individual level. We don't believe that technology should be stifled just in order to satisfy copyright holders. Technology today has already been stifled because of concerns about copyright issues.

CPU: How so?

Bildson: For example, up until last October, we generated no revenue. Part of the reason that we couldn't raise any outside money was the whole Internet bubble burst. But certainly there would be a thousand file sharing and P2P search companies available in the market right now, probably funded by various VCs and what-not, if it wasn't for this fear that they were going to get sued out of existence. Another valid example would be JXTA technology from Sun [which lets any connected device on a network communicate in a P2P fashion]. JXTA was extremely interested in having someone like LimeWire provide

search-and-sharing capabilities because they were scared to provide them themselves due to legal issues.

CPU: Do you feel that the legal issue would disappear if digital rights management were universally and adequately implemented on the files?

Bildson: We would definitely favor digital rights management being used. Someday you're going to have a marketplace where you can download DRM-protected material. You'll probably be able to sample a little bit of it and then purchase the file. There's going to be free content that uses digital rights management to lock in commercial material with it, like the J!ve Player linked on our home page does. The J!ve Player builds in shopping features and advertising into the content they distribute. We see both of those as extremely legitimate businesses that are going to be helped immensely by file-sharing capabilities.

CPU: Is this the future: Say Aerosmith releases a new song wrapped up in a tight DRM package. I go out and download it from the publisher, and I pay them \$5 for the song, so I have the rights to use this file however the rights specify. Now, I can make this file available for sharing via LimeWire over some other P2P network. If somebody downloads the song from me, then when they try to play it, up pops a window that says, "Do you want to purchase the rights for this? Give me your \$5."

Bildson: Yes. Once DRM-locked files became available, they would quickly be spread on the file-sharing networks; then you could unlock them locally with a key. Video is a great example, where once there is broadband, you can download a DVD-quality movie. If you want to watch the movie, you pay some money to unlock it and there you go. It just makes sense to me that this is going to be the way of the future. **CPU**

To read our entire interview with Greg Bildson, go to
www.smartcomputing.com/cpumag/jun02/bildson

Under Development

A Peek At What's Brewing In The Laboratory

Fresh from the most influential R&D labs around the world, here's a glimpse at some of the technology that scientists, lab techs, and researchers are cooking up for the future.

Mind vs. Mouse

Since the dawn of computing, writers have dreamed of being able to think their words onto the screen (this writer has, anyway), and gamers have pined to mentally control their on-screen personae. After all, the fraction of a second it takes for an impulse to travel from brain to finger could be the difference between life and game over. Research in the '90s into melding mind and PC by way of implanted electrodes showed marginal success at best. Some subjects could think a cursor into halting movement; other subjects found the learning curve too steep and failed entirely. Now, however, Brown University researchers have hit on a technique so successful, even a monkey can do it.

Using an implant chip roughly the size of a shirt button, lead researcher Dr. Mijail Serruya and his biomedical team have taken a monkey and wired its brain with an array of 100 electrodes connected to the tiny controller. Each electrode is connected to a neuron and senses electrical impulses crossing those cells. The activity is relayed to the controller, which is wired to a PC. Software running a specialized algorithm translates the impulse activity into cursor movement on a monitor screen.

Serruya has experimented with three monkeys so far, and the results are phenomenal. A monkey's task is to direct the cursor mentally and chase a moving target around the screen. According to the Brown researchers, the simians could

"instantly" control the interface and control the cursor with a speed and accuracy that rivaled their performance when using a mouse.

Although no one is seriously discussing the use of brain implants for a rousing game of Halo, the technology is tantalizingly close to offering hope for paralysis victims. Whereas several of today's interfaces for the physically challenged involve translating head or eye movements into cursor commands,

Brown University's Dr. Mijail Serruya is using brain implants and a special new processing system to make mental control of an on-screen cursor so easy, even a monkey can do it.



brain control would likely offer far greater performance, letting users do everything from communicating by e-mail to carrying out Web-enabled occupations, such as writing or research.

The Brown team's efforts have shown such promise that it has filed for a patent on the approach and founded a company named Cyberkinetics to work on developing medical applications. Before you stand in line for your wetware-based mouse replacement, though, keep in mind that the technology is likely at least a decade away from delivering a commercial product. ▲

Read My Lips

Many posh restaurants around the country have a "no cell phone" policy, and taking a call in the middle of a movie will likely get you doused in Coke and popcorn. However, NTT DoCoMo (www.nttdocomo.com), Japan's leading telecom provider, is working on a technology that will let you carry on a conversation in absolute silence.

According to a *New Scientist* report, DoCoMo is developing a handset that uses a contact sensor pressed against the skin by one's mouth to detect electrical signals flowing through nearby muscles. The different muscle movements involved in speaking produce different signal patterns, which can be translated into a synthesized voice or text—an intriguing prospect for hands-free instant messaging.

Today's accuracy is only sufficient for vowel recognition, but researchers are working on mastering consonants. Handset-mounted cameras, a soon-to-be-common feature in 3G phones, may one day watch the user's mouth to enhance accuracy.

Although DoCoMo's research aims primarily to reduce public noise—a serious problem in overcrowded urban Japan—the technology would also enable a way for the mute to communicate verbally. DoCoMo anticipates that the technology should be ready to market in roughly five years. ▲



NTT DoCoMo is working on a system that senses muscle movements around the mouth, turning a cell phone user's lip motions into text or speech, without making any noise at all.

A Keyboard You Can Wear

If you thought Think Outside's (www.thinkoutside.com) folding keyboard was nifty for typing on the road, get ready for something that's even more flexible—literally.

British switching and sensing company ElekSen (www.eleksen.com) has hit the market with its first product based on the company's patented ElekTex technology. ElekTex is a durable, water resistant material interwoven with flexible, conducting fibers. The fabric need only be 1mm thick to sense pressure along its X, Y, and Z axes.

The first example of ElekTex is Logitech's (www.logitech.com) new Key-Case. Lightweight and spillproof, the KeyCase is a thin protective fabric shell that wraps around your Palm PDA. Unfold the case, however, and the shell's interior unfolds to reveal a flat keyboard. Naturally,

the design takes a little getting used to. Keys lack the tactile feedback of a traditional keyboard, and because the Key-Case is flexible, you should use it on a hard, flat surface. ElekTex technology is surprisingly cheap, selling for only \$99.95, which is approximately the same price as a Think Outside Stowaway Portable Keyboard.

ElekTex products can be designed as either self-contained peripherals, such as the KeyCase, or interface add-ons, such as a touch pad possibly integrated into the face of a steering wheel. A soft cell phone is in the works that you could twist, dunk, and drop without fear of shattering. Future applications may include everything from gaming controllers to a multifunction watch that wraps around your wrist. ▲

PlayStation Or Workstation?

Roughly a year ago, Todd Martinez, a professor in the Department of Chemistry at the University of Illinois at Urbana-Champaign, purchased a Sony PlayStation 2. When he asked the business office to reimburse him, officials were less than kind.

"Accounting started yelling at me. They didn't want to reimburse me, but they finally did when my department head stepped in and explained that this was for real research. When the inventory person came to put a tag on it, she was laughing for a week."

Martinez and his team aren't simply looking to play a few rounds of Final Fantasy after their lab work is done, though. The PlayStation 2 is the lab work.


Although the gaming console only runs on a 300MHz CPU, the chip is capable of parallel processing 10 floating-point pipelines. The linear algebra equations that fuel cutting-edge gaming graphics are also at the core of quantum chemistry. This esoteric field aims to solve equations that reveal how electrons are distributed within the molecules of materials. Figure this out and you can predict how those molecules and materials will behave under different circumstances. Anything from designing drugs to optical switching tools to smaller features on

computer chips can apply quantum chemistry. All it takes is a lot of processing horsepower to solve the equations.

"What really interested us in the PlayStation 2 was that it has a theoretical gigaflop rating six times greater than a 1GHz Pentium III," says Martinez. "Sure, it's easier to develop for workstations than a PlayStation, but the PlayStation is more cost efficient. Because the price is so much lower, you can throw three or four times as many machines at the problem."

Martinez purchased one of the only 2,000 Linux development kits that Sony released in Japan. Having to construct PS-compatible Linux apps from scratch has been a time-consuming challenge, although the group is constantly figuring out how to improve the code and achieve greater performance. Martinez expects that other universities will soon build on the work his group is pioneering. Ultimately, conducting research with off-the-shelf consumer equipment such as the PlayStation may make all the difference in research groups being able to carry out their projects, especially perennially tight-fisted academic circles. ▲

Your Eyes Say You're Lying



Imagine standing in line at an airport check-in. An attendant next to a camera points at the face of a man in front of you and asks him if he's carrying any sharp objects or explosives. He says, "No," even though he has a pocketknife hidden in his jeans. Seconds later, guards swoosh out of nowhere, leading him away for questioning.

Especially now, with facial recognition and voice analysis both under fire for their inaccuracy in assessing criminal identity or intent, a new technology developed by Dr. Ioannis Pavlidis, Honeywell Laboratories senior principal research scientist, and Dr. James Levine of the Mayo Clinic, may help pick out liars with little more than a simple question and answer. The technique hinges on the fact that people tend to experience increased blood flow and warming around their eyes when lying. An infrared camera captures a "thermal signature" of a person's face in a normal state, and then looks again as the person replies to a question.

In testing by the U.S. Department of Defense, a standard polygraph machine detected 70% of all liars. Honeywell's thermal imaging system snagged 83%. Although the system was in development about a year before Sept. 11, the technology is being fast-tracked and developed in cooperation with multiple government agencies, including the Defense Advanced Research Projects Agency. ▲


Back Door

Q&A With Federico Faggin

The phrase “creative genius” becomes something of a dry understatement when used to describe Dr. Federico Faggin, co-inventor of the microprocessor. The modern microchip made possible the personal computer revolution of the 1970s and 1980s. And the former Intel insider, who is now celebrating the 30th anniversary of his computing creation, is still going strong. His company, San Jose, Calif.-based Synaptics, makes touch-screen technologies and Chinese-language handwriting-recognition systems. In late January 2002, the company went public on NASDAQ. We recently interviewed Dr. Faggin; what follows are the highlights of that conversation.

Q The anniversary of the microprocessor has arrived. Tell us about how you got involved in that project.

FAGGIN: This was the Intel 4004 microprocessor. It was 30 years ago. I got involved because I was hired by Intel. I had a contract. I joined Intel in April 1970, and the story of the 4004 goes back to about a year earlier. The Japanese calculator maker *Busicom* came to Intel with seven custom chips that they wanted Intel to make. At that time, the application manager took a look at the project and he wanted a programmable system. It was really state-of-the-art for what people were doing for calculators in those days. They were using some level of programmability. There were macro instructions. Shift registers were the memory. RAM was only beginning to be done with memory. It was a special purpose architecture. So we developed an architecture that was more like that for a computer, an instruction set that made sense. People at that point had been talking about making a CPU on a chip



for some time, but because of silicon gate technology, it was doable.

Q So the applications department assigned you as the project leader?

FAGGIN: Yes. There were four chips. One was a CPU. One was a ROM with I/O [input/output]. Another was a RAM with I/O. Another one was a static shift register for I/O applications. So basically, with help from the engineer at *Busicom*, who helped with the definition of the technology, I ended up designing the chip set. I developed all the methodology that was necessary to make random logic with the silicon gate. Intel was a memory company. They had never done any random logic at that point. That was a major task to be able to do that.

Q So you took over from a set of customer specifications?

FAGGIN: Yes. That's it. The project was the chipset for a number of calculators and calculator-like machines, like adding machines and cash

registers. Calculator-based equipment, basically. The first project, though, was a four-function, printing calculator with memory. It was a desktop machine.

Q Did you think it could have multiple uses, the chipset, that is?

FAGGIN: We did. In October of 1969, another company went to Intel. They were called CTC. Later on, they changed their name to *DataPoint*. They were a prominent company in the 1970s and 1980s in intelligent terminals. They had developed an intelligent terminal and there was a little computer there too at the heart of it. They wanted Intel to do a special memory for them. They needed an odd-sized memory. Intel took a look at this. They had just finished the architecture for the 4004. The architecture required for the CTC computer was very similar, except this was a 9-bit, RAM-based machine, rather than a ROM-based machine, like the *Busicom* machine. Intel said that they could probably do it all on one chip for CTC, which was a gutsy thing to say, because the 4004 was still not done yet. They were banking on someone in-house being able to do it. Based on that, CTC entered into a contract with Intel. There was another project: the 1201, which later became the 8008, the first independent microprocessor that Intel made.

For our complete interview with Federico Faggin, go to www.smartcomputing.com/cpumag/jun02/faggin.

Mr. Koprowski has been a business journalist for nearly two decades, covering computing and technology for PBS, CBS, and The Wall Street Journal Europe and WSJ Interactive Edition.





check is useless for things words can't describe.

While we can't put into words the impact of a Nikon digital image, we can tell you a little about the camera that created this one. Introducing the Coolpix® 995. It comes fully loaded with 3.34 megapixel resolution, 4x Optical Zoom-Nikkor lens, automatic exposure with manual  override, a pop-up Speedlight, and a quick review feature that instantly allows you to scroll through your images. With this many options, the only limit is your imagination. Visit nikonusa.com or call 1-800-NIKON-UX.



The Coolpix 995.

WHAT'S YOUR BURNING DESIRE?

VIDEO 

AUDIO 

IMAGES 

DATA 

Whatever it is, chances are the Panasonic LF-D321U will satisfy it. That's because it burns DVDs using only DVD Forum approved formats like DVD-RAM and

DVDBurner
DVD-RAM/R DRIVE

DVD-R. Only DVD Forum approved formats

ensure information security and maximum compatibility between DVD players.

On the Read side, it's just as flexible, because it reads virtually every DVD and CD format. So you get extreme compatibility with extreme storage capacity—up to 9.4 GB per disk!

It's perfect for Video, Data and Imaging applications.



Also included is a set of powerful software titles that gives you the ability to play, capture, edit and burn your Video files, as well as back-up your hard drive.



So if you're looking for the hottest DVD Drive that gives you the power to burn and more, check out the Panasonic LF-D321U today.

For more information, visit
www.panasonic.com/storage

Panasonic



DVD-CPU